Draft version 25 August 2025
Typeset using LaTeX **twocolumn** style in AASTeX631

GWTC-4.0: Population Properties of Merging Compact Binaries

The LIGO Scientific Collaboration, The Virgo Collaboration, and The KAGRA Collaboration

(Compiled: 25 August 2025)

ABSTRACT

We detail the population properties of merging compact objects using 158 mergers from the cumulative Gravitational-Wave Transient Catalog 4.0, which includes three types of binary mergers: binary neutron star, neutron star–black hole binary, and binary black hole mergers. We resolve multiple over- and under-densities in the black hole mass distribution: features persist at primary masses of $10\,M_\odot$ and $35\,M_\odot$ with a possible third feature at $\sim 20\,M_\odot$. These are departures from an otherwise power-law-like continuum that steepens above $35\,M_\odot$. Binary black holes with primary masses near $10\,M_\odot$ are more likely to have less massive secondaries, with a mass ratio distribution peaking at $q = 0.74^{+0.13}_{-0.13}$, potentially a signature of stable mass transfer during binary evolution. Black hole spins are inferred to be non-extremal, with 90% of black holes having $\chi < 0.57$, and preferentially aligned with binary orbits, implying many merging binaries form in isolation. However, we find a significant fraction, 0.24–0.42, of binaries have negative effective inspiral spins, suggesting many could be formed dynamically in gas-free environments. We find evidence for correlation between effective inspiral spin and mass ratio, though it is unclear if this is driven by variation in the mode of the distribution or the width. The binary black hole merger rate increases with redshift as $(1+z)^\kappa$ with $\kappa = 3.2^{+0.94}_{-1.00}$, consistent with the cosmic star formation density. While there is no evidence of the mass spectrum evolving with redshift, the distribution of effective inspiral spin is found to broaden as redshift increases out to $z \approx 1$. We infer the local merger rates (i.e., at redshift $z = 0$) to be 7.6–250 $\mathrm{Gpc^{-3}\,yr^{-1}}$ for binary neutron stars, 9.1–84 $\mathrm{Gpc^{-3}\,yr^{-1}}$ for neutron star–black hole binaries, and 14–26 $\mathrm{Gpc^{-3}\,yr^{-1}}$ for binary black holes; all values reflect central 90% credible intervals.

1. INTRODUCTION

Gravitational waves (GWs) have revolutionized the study of compact objects and their populations (Abbott et al. 2016, 2019a, 2021a, 2023a), with observations now extending beyond redshift $z = 1$. Searches for GWs from compact binary mergers have robustly quantifiable selection effects, which enable detailed inference of population properties with few sources of potential bias (Mandel et al. 2019; Essick & Fishbach 2024). By studying the growing catalog of compact binary mergers, we aim to uncover both where these systems form, and how they evolve toward merger. For example, are they formed from stars that are gravitationally bound at birth, possibly as the products of chemically homogeneous evolution (Mandel & de Mink 2016; Marchant et al. 2016; de Mink & Mandel 2016), or brought close enough together for GW-driven inspiral through phases of common envelope evolution (Bethe & Brown 1998; Portegies Zwart & Yungelson 1998; Belczynski et al. 2001; Dominik et al. 2015) or stable mass transfer (Hurley et al. 2002; Neijssel et al. 2019; van Son et al. 2022b)? Could binaries be assisted in assembly and hardening in the gaseous disks of active galactic nuclei (AGN) (McKernan et al. 2012; Bartos et al. 2017; Stone et al. 2017; Fragione et al. 2019) or through dynamical interactions in dense stellar environments (Kulkarni et al. 1993; Sigurdsson & Hernquist 1993; Portegies Zwart & McMillan 2000; Ziosi et al. 2014)?

From the first three observing runs of the LIGO–Virgo–KAGRA Collaboration (LVK), we established that stellar-mass black holes (BHs) in merging binaries have a broad distribution of masses, with peaks at primary masses of $\sim 10\,M_\odot$ and $\sim 35\,M_\odot$, that falls off steeply above $\sim 45\,M_\odot$ (Abbott et al. 2023a). Despite the steep decline, the rate of mergers in the expected pair-instability supernova (PISN) gap, predicted to begin at $\sim 45 - 50\,M_\odot$ (Woosley 2017; Farmer et al. 2019), was found to be non-zero (Abbott et al. 2020c,d). Likewise, the rate of mergers in the purported lower mass gap between $\sim 3 - 5\,M_\odot$ (Bailyn et al. 1998; Ozel et al. 2010; Farr et al. 2011) was found to be small but non-zero (Abbott et al. 2020e; Abac et al. 2024). The neutron star (NS) mass distribution does not require a peak in the distribution at $\sim 1.33\,M_\odot$, as is seen in

the galactic binary neutron star (BNS) population (Özel & Freire 2016; Farrow et al. 2019). The binary black hole (BBH) merger rate was found to definitively increase with redshift. Spins were found to be small in magnitude, with a non-zero fraction of systems with spin components anti-aligned with the binary orbit.

Advanced LIGO (Aasi et al. 2015), Advanced Virgo (Acernese et al. 2015), and KAGRA (Akutsu et al. 2021) began their fourth observing run (O4) on 2023 May 24 at 15:00 UTC. The first part of the fourth observing run (O4a) ended on 2024 January 16 at 16:00 UTC. The accompanying Gravitational-Wave Transient Catalog (GWTC) version 4.0 (hereafter, GWTC-4.0) (Abac et al. 2025a,b,c) contains those events observed in previous observing runs, O1 (Abbott et al. 2016), O2 (Abbott et al. 2019b), and O3 (Abbott et al. 2021b, 2023b, 2024), together with the newest observations from O4a. We use the updated catalog to infer the population properties of BNS, neutron star–black hole binary (NSBH), and BBH systems in the local Universe.

GWTC-3.0 included 76 candidates with false alarm rate (FAR) $< 1\,\mathrm{yr}^{-1}$: 69 BBHs, 4 NSBHs, 2 BNSs, and one event GW190814_211039 (henceforth, GW190814) that is either a NSBH or BBH. Abac et al. (2025c) identifies 128 candidate signals in O4a with a probability of astrophysical compact binary origin of $p_{\mathrm{astro}} \geq 0.5$, of which 85 (84 BBHs and 1 NSBH) have FAR $< 1\,\mathrm{yr}^{-1}$. Selecting BBHs from the catalog using this threshold yields a cumulative BBH count of 153. With fewer signal candidates from binaries containing at least one NS, maintaining a comparable contamination fraction to that of BBH mergers ($\sim$5%) requires a more conservative threshold (see Section 3 for more details). We adopt the same threshold used for similar analyses of GWTC-3.0, FAR $< 0.25\,\mathrm{yr}^{-1}$, for analyses that include NS-containing populations (Abbott et al. 2023a). This yields two NSBH candidates from O3, and one new candidate, GW230529_181500 (henceforth, GW230529), detected above this threshold in designated observing time during O4a (Abac et al. 2024). Section 3 provides further details and discussion of threshold choices and candidate inclusion. Adopting the more conservative FAR $< 0.25\,\mathrm{yr}^{-1}$ threshold reduces the O4a BBH count to 76, resulting in a catalog of 2 BNSs, 3 NSBHs, and 138 BBHs for analyses that include BNS and NSBH populations in this work.

The remainder of the paper is structured as follows. In Section 2 we provide a brief description of our inference techniques (with remaining details in Appendix A) and the classes of models used (with full descriptions of the models in Appendices B and C, and how they were chosen in Appendix D). Section 3 describes our dataset and sample selection, including brief descriptions of search techniques, threshold choices, and waveform models used. In Section 4 we present measurements of the complete compact-binary mass spectrum (NSs and BHs). In Section 5, we study the properties of the NS-containing population in detail, and in Section 6 focus on the ensemble properties of BBHs, including their masses, spins, redshifts, and associated correlations. Section 7 concludes by summarizing the key results of this work. Associated data releases provide analysis results and figure generation scripts (LIGO Scientific Collaboration et al. 2025) and data products for estimating sensitivity (Essick 2025a,b).

2. METHODS

We determine the population properties of merging compact binaries using hierarchical Bayesian inference as has been done in previous studies (Abbott et al. 2019a, 2021a, 2023a). Our aim here is to estimate the posterior distribution $p(\boldsymbol{\Lambda}|d)$ on population-level model parameters $\boldsymbol{\Lambda}$ (also referred to as *hyperparameters*) in light of new data d representing the observed GW data from individual merger events. These model-dependent hyperparameters describe the population-level properties of GW source parameters $\boldsymbol{\theta}$ (masses, spins, redshifts, etc.). According to Bayes' theorem,

$$p(\boldsymbol{\Lambda}|d) = \frac{\mathcal{L}(d|\boldsymbol{\Lambda})\pi(\boldsymbol{\Lambda})}{p(d)}\,, \tag{1}$$

where $\pi(\boldsymbol{\Lambda})$ is the prior probability distribution and $\mathcal{L}(d|\boldsymbol{\Lambda})$ is the likelihood—the probability of obtaining the data d given some $\boldsymbol{\Lambda}$. The Bayesian *evidence* $p(d)$ ensures that $p(\boldsymbol{\Lambda}|d)$ is properly normalized.

In order to use Bayes' theorem to infer the posterior probability distribution of hyperparameters, we need the likelihood of obtaining the observed catalog of events given a set of hyperparameters and a population model. Because GW detectors are not equally sensitive to different astrophysical sources, the likelihood must account for selection biases. Assuming (i) a Poisson process generates realizations from the source population of which a subset is detected, and (ii) the observed data associated are statistically independent (e.g., no overlapping signals), the likelihood is given by (Loredo 2004; Farr et al. 2015; Mandel et al. 2019; Thrane & Talbot 2019; Vitale et al. 2020)

$$\mathcal{L}(\{d_i\}, N_{\mathrm{det}}|\boldsymbol{\Lambda}) \propto N(\boldsymbol{\Lambda})^{N_{\mathrm{det}}} e^{-N_{\mathrm{exp}}(\boldsymbol{\Lambda})} \prod_{i=1}^{N_{\mathrm{det}}} \int \mathrm{d}\boldsymbol{\theta}\, \mathcal{L}(d_i|\boldsymbol{\theta})\pi(\boldsymbol{\theta}|\boldsymbol{\Lambda})\,, \tag{2}$$

where $N(\mathbf{\Lambda})$ is the total number of mergers (detected and undetected) incident on the detectors within the observing period, $N_{\rm det}$ is the number of detections, d_i are the strain data corresponding to the ith detection, and $\mathcal{L}(d_i|\boldsymbol{\theta})$ is the likelihood of the data d_i given the GW source parameters $\boldsymbol{\theta}$ (Abac et al. 2025b). The expected number of detections is $N_{\rm exp}(\mathbf{\Lambda}) = N\,\xi(\mathbf{\Lambda})$, with $\xi(\mathbf{\Lambda})$ being the expected fraction of the population parameterized by $\mathbf{\Lambda}$ that is detectable. Formally, we impose a detection criterion on the data, e.g., a FAR threshold, by the selection function $x(d)$ acting on a data segment d. A GW event is detectable in d if $x(d) > x_{\rm thr}$, where $x_{\rm thr}$ is the chosen detectability threshold. Then,

$$\xi(\mathbf{\Lambda}) = \int_{x(d)>x_{\rm thr}} \mathrm{d}d\,\mathrm{d}\boldsymbol{\theta}\, p(d|\boldsymbol{\theta})\pi(\boldsymbol{\theta}|\mathbf{\Lambda}). \tag{3}$$

In Equation (3), the domain of integration is over all data d which surpasses the detection threshold $x_{\rm thr}$ (Essick et al. 2025). When a prior $\pi(N) \propto 1/N$ is assumed, Equation (2) can be analytically marginalized over N leaving the rate-marginalized hierarchical likelihood

$$\mathcal{L}(\{d\}, N_{\rm det}|\mathbf{\Lambda}) \propto \prod_{i=1}^{N_{\rm det}} \frac{\int \mathrm{d}\boldsymbol{\theta}\,\mathcal{L}(d_i|\boldsymbol{\theta})\pi(\boldsymbol{\theta}|\mathbf{\Lambda})}{\xi(\mathbf{\Lambda})}\,. \tag{4}$$

This likelihood may also be obtained by marginalizing over a prior on the expected number of detections, $N_{\rm exp} = N\xi(\mathbf{\Lambda})$, rather than on N (Essick & Fishbach 2024).

Having described the hierarchical likelihood in Equation (2) and Equation (4) we have one of the two ingredients necessary for sampling the posterior in Equation (1). We also must choose a prior over the space of hyperparameters. This depends on the population model, which we describe in more detail below. We specify the priors for each model in Appendix B.

Equation (2) is the exact form of the likelihood. However, calculating $\xi(\mathbf{\Lambda})$ and integrating over $\boldsymbol{\theta}$ is analytically intractable. Therefore, we approximate the likelihood via Monte Carlo reweighting using importance sampling. Our Monte Carlo approximation for the likelihood carries some uncertainty and may not be converged appropriately for some hyperparameter values. Therefore, we discard hyperparameters whose likelihood uncertainty exceeds a chosen threshold *a posteriori* (Talbot & Golomb 2023). For more description of this problem and our approach for mitigating it, see Appendix A.

The Bayesian inference problem is stated here in terms of the population probability density $\pi(\boldsymbol{\theta}|\mathbf{\Lambda})$ and the overall number of merging binaries N. This can be converted to another astrophysical quantity of interest, the comoving source-frame merger rate density

$$\mathcal{R}(z) = \frac{\mathrm{d}N}{\mathrm{d}V_{\rm c}\mathrm{d}t_{\rm s}}(z) = \frac{\mathrm{d}N}{\mathrm{d}t_{\rm d}\mathrm{d}z}\left(\frac{\mathrm{d}V_{\rm c}}{\mathrm{d}z}\frac{1}{1+z}\right)^{-1}, \tag{5}$$

where $t_{\rm s}$ is the time measured in the comoving source frame, $t_{\rm d}$ is the time at the detector (redshift $z = 0$) and $\mathrm{d}V_{\rm c}/\mathrm{d}z$ is the differential comoving volume with respect to redshift z (see e.g., Essick et al. 2025). The comoving source frame merger rate represents the number of mergers in a unit of comoving volume and source-frame time, conventionally measured in units $\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$.

To construct models for the population distribution of astrophysical GW sources, we take one of two approaches. The first approach—which we call the *strongly modeled approach* (sometimes called the parametric approach)—assumes a specific functional form $\pi(\boldsymbol{\theta}|\mathbf{\Lambda})$ for the astrophysical distribution *a priori*, e.g., a Gaussian distribution or a power law. A second approach—which we call the *weakly modeled approach* (elsewhere data-driven, flexible, or nonparametric)—attempts to make minimal *a priori* assumptions about the underlying astrophysical population, e.g., a spline model. We elaborate on each approach in the following two subsections, and provide details for the models presented in this work falling in each category.

2.1. *Strongly Modeled Approach*

The strongly modeled approach assumes that the underlying astrophysical distribution of binary properties can be described by a fixed functional form and associated hyperparameters. Consequently, this approach has far fewer hyperparameters as compared to the weakly modeled approach. The parameterization may be motivated by theoretical expectations about the astrophysical population and/or could be selected because it seems to fit the observed data well. This has the benefit of being simple and interpretable, as hyperparameters can be designed to correspond directly to physical features of interest. Examples of such features are a distribution's minimum or maximum, the location parameter for an overdensity corresponding to e.g., pulsational pair-instability supernovae, etc. On the other hand, these approaches can be overly restrictive: features present in the true astrophysical distribution that are not captured by our parameterization cannot be easily discovered. Additionally, the inferred parameters can be biased due to a mis-specified model (e.g., Romero-Shaw et al. 2022).

The parameterized models selected for this work are generally simple extensions to those employed in the previous LVK astrophysical population studies paper from

Table 1. Summary of Models

Model Type	Model Name	Description
Strongly Modeled: Mass	FullPop-4.0	Models the mass spectrum of all CBCs simultaneously with appropriate power-law and peak components. Also allows for a gap between the most massive NS and the least massive BH.
	Broken Power Law + 2 Peaks ★	The primary mass distribution has a broken power law continuum between a minimum and maximum mass, plus two Gaussian peaks around $\sim 10\,M_\odot$ and $\sim 35\,M_\odot$. The distribution of mass ratio q is a power law between some minimum value and 1.
	Extended Broken Power Law + 2 Peaks	The mass-ratio power-law is allowed to differ between primary masses in the continuum and in the $35\,M_\odot$ peak.
Strongly Modeled: Component Spin	Gaussian Component Spins ★	The spin magnitude population is a Gaussian truncated over the physical range $\chi \in [0, 1)$. The distribution of cosine spin tilts relative to the orbital angular momentum includes an isotropic (i.e., uniform) component and a truncated Gaussian component.
Strongly Modeled: Effective Spin	Gaussian Effective Spins	The joint $\chi_{\rm eff}$–$\chi_{\rm p}$ effective spin distribution is a bivariate Gaussian allowing for correlations.
	Skew-normal Effective Spin	The $\chi_{\rm eff}$ marginal effective spin distribution is skew normal, truncated to $[-1, 1]$.
Strongly Modeled: Redshift	Power Law Redshift ★	The merger rate per unit comoving volume and source-frame time evolves with redshift z as a power law i.e., $\propto (1+z)^\kappa$.
Strongly Modeled: Correlations	Copula	Truncated Gaussian distributions are assumed for $\chi_{\rm eff}$ and $\chi_{\rm p}$. A Frank copula density function correlates two variables. Separate copula models correlate $(q, \chi_{\rm eff})$, $(z, \chi_{\rm eff})$, $(m_1, \chi_{\rm eff})$, and (m_1, z).
	Linear	A truncated Gaussian distribution is assumed for $\chi_{\rm eff}$ with the mean and width linearly dependent on either q or z.
	Spline	A truncated Gaussian distribution is assumed for $\chi_{\rm eff}$ with the mean and width dependent on either q or z. This dependence is flexibly modeled with a cubic spline.
Weakly Modeled: All Parameters	B-Spline	Fits the astrophysical distribution as a separable joint distribution with one-dimensional basis splines. Large numbers of basis functions allow for flexibility, with difference-based priors imposing smooth evolution *a priori*.
	Binned Gaussian Process	Assumes a fixed binning scheme and infers the event rate under the assumption of a constant rate within each bin, and a Gaussian process prior imposing smooth covariance across bins.

Strongly modeled and weakly modeled approaches used to study the mass, spin, and redshift distributions of merging compact binaries. We provide a brief description of each model here, with detailed descriptions in Appendix B and Appendix C. Models marked with a ★ are treated as defaults and are used whenever no model is explicitly indicated for a certain parameter. All strongly modeled approaches mentioned below target BBHs, with the exception of FullPop-4.0 that is used to model the entire CBC population.

GWTC-3.0 (Abbott et al. 2023a). Updates to these models are motivated by features that have emerged due to new data in GWTC-4.0 and improved interpretations since GWTC-3.0.

Table 1 lists the strongly modeled approaches used in this paper, with details in Appendix B. We also describe the parameterizations and priors for these models in Appendix B, and our procedure for selecting the default strongly modeled approach in Appendix D. For most models, we assume that masses, spins, and redshifts are all uncorrelated. We also study a selection of pairwise correlations between parameters in Section 6.5.

2.2. *Weakly Modeled Approach*

The weakly modeled approach adopts models that deliberately make few assumptions about the nature of the underlying compact-binary population. Such approaches typically require a larger number of hyperparameters in order to effectively approximate a wide variety of distributions. The philosophy of a weakly modeled approach is to discover unexpected features in the population, which may be unforeseen or difficult to parameterize. However, they could yield results that are more difficult to interpret astrophysically. The difference between the strongly modeled and weakly modeled approaches can be understood in terms of the bias–variance tradeoff; the former has low variance with a risk of bias, whereas the latter has low bias but elevated variance.

Weakly modeled approaches must make some assumptions, however, and must be designed with different features in mind. For example, different approaches have sensitivity to astrophysical correlations, narrow structures, or gaps in the population. While a unified Bayesian approach to capture generic population features is still an open problem (Mandel et al. 2017; Tiwari 2021; Rinaldi & Del Pozzo 2021; Edelman et al. 2023; Golomb & Talbot 2023; Payne & Thrane 2023; Toubiana et al. 2023; Callister & Farr 2024; Ray et al. 2023a; Farah et al. 2025b; Heinzel et al. 2025a), we use two different weakly modeled approaches—B-Spline and Binned Gaussian Process (BGP)—to verify results from our strongly modeled approaches (see Table 1). We describe two additional weakly modeled approaches—Autoregressive Process (AR) and Flexible Mixtures (FM)—in Appendix C, and compare these approaches in Appendix D.4.

3. DATASET

3.1. *Data Collection Duration*

Analyses presented in this paper use selected data products from GWTC-4.0 (Abac et al. 2025a,b,c,d). This section and Section 3.2 provide details on the selection criteria for events analyzed for this paper. GWTC-4.0 includes GW candidates and data from O1 through the end of O4a, as well as a GW candidate and data collected during an engineering run (ER) (Abbott et al. 2020a) directly preceding the start of O4a. ERs are periods dedicated to final commissioning and configuration of the instruments prior to an observing run; the instruments may be in locked and low-noise configurations, but are not generally intended to perform astrophysical observations. The ER data included in GWTC-4.0 is deliberately chosen to contain a few days of data around a GW event potentially originating from a NSBH binary merger, GW230518_125908 (henceforth, GW230518). The analyses and results quoted in this paper exclude data obtained during the ER; the inclusion of this data would introduce human selection effects that cannot be easily incorporated into $\xi(\mathbf{\Lambda})$, and hence may bias our inferences.

3.2. *Event Selection Criteria*

3.2.1. *Significance Thresholds*

To ensure that the dataset we use in this paper has reduced contamination from noise events, we adopt a significance threshold of FAR < 1 yr^{-1} in at least one GW search pipeline, which is consistent with the criterion adopted in Abbott et al. (2023a). Based on the FAR threshold, a total of 161 CBC candidates have been detected from O1 through O4a (Abbott et al. 2019b, 2021b, 2024, 2023b) by the GW search pipelines (Abac et al. 2025c), of which 85 were from O4a. This is a noteworthy increase in the number of observations reported in Abbott et al. (2023b), which contained 76 events meeting the FAR < 1 yr^{-1} threshold. With this FAR threshold and assuming noise signals are produced independently, we expect $\sum_k \mathrm{FAR} \times T_k \simeq 6.7$ contaminant noise events in our results, where T_k is an estimate of the time examined by the kth search. The list of GW events included in the analyses of this paper contains GW231123_135430 (henceforth, GW231123), which has high probability for being the most massive BBH with FAR < 1 yr^{-1} detected to date by the GWs with both component masses possibly in the upper mass gap (Abac et al. 2025e; Woosley 2017; Mapelli et al. 2020; Farmer et al. 2019, 2020; Woosley & Heger 2021; Hendriks et al. 2023). Not all events reported in the GWTCs papers (Abbott et al. 2019b, 2021b, 2024, 2023b; Abac et al. 2025c) are included in our analyses, as previous GWTC papers thresholded event candidates using $p_{\mathrm{astro}} \geq 0.5$ or FAR < 2 yr^{-1}, whereas the analyses presented in our paper select events with a significance of FAR < 1 yr^{-1}. Here, p_{astro} is the estimate of the probability of

astrophysical origin of the event candidates (Abac et al. 2025b,c). In addition, GW230630_070659 is excluded from the analyses and the number of events reported in this paper, due to concerns of the data quality around the time of this event (Abac et al. 2025c).

3.2.2. *Mass and Significance Thresholds for Events with NSs*

To distinguish NS-containing events from events containing only BHs, we first threshold events by checking whether the 1% lower limit on the component mass is smaller or larger than $3\,M_\odot$. As far fewer GW candidates with NSs have been observed compared to BBHs, this paper adopts a stricter FAR threshold of < 0.25 yr^{-1} in at least one GW search pipeline for GW candidates with NSs to ensure a purer sample, as done in Abbott et al. (2023a). This FAR threshold excludes GW190917_114630 and GW190426_152155 (Abbott et al. 2024), which are consistent with originating from NSBHs but have FARs of > 0.25 yr^{-1}. In addition, we exclude certain events whose category is ambiguous from dedicated BBH and NSBH analyses. Specifically, we exclude GW190814 as its source's secondary mass is lower than the component masses of events classified as BBHs but higher than the inferred NS mass range, leaving its classification ambiguous (Abbott et al. 2020e; Essick et al. 2022; Abbott et al. 2023a). This event is included in Section 4 which considers binary-merger populations across all masses. Hence, the only NS-containing event detected in O4a considered in this paper is GW230529 (Abac et al. 2024), in addition to the previously reported NSBHs, GW200105_162426 and GW200115_042309 (henceforth, GW200105 and GW200115, respectively). The results are presented in Section 5.

3.2.3. *Exclusion of Non-LVK Catalog Events*

In addition to the catalog of event candidates and its analyses conducted by the LVK, independent teams have analyzed the public GW data from O1 through second half of the third observing run (O3b) using alternative algorithms and have identified additional GW binary-merger event candidates (Venumadhav et al. 2019, 2020; Olsen et al. 2022; Mehta et al. 2025; Wadekar et al. 2023; Zackay et al. 2019; Nitz et al. 2019, 2020, 2021, 2023; Kumar & Dent 2024; Mishra et al. 2025; Koloniari et al. 2025). We do not include these additional events in our analyses here due to subtleties with consistently combining sensitivity estimates from these independent catalogs.

3.3. *Sensitivity of GW Searches*

A key ingredient in the estimation of population level properties is the sensitivity of our GW searches $\xi(\mathbf{\Lambda})$. The following four search pipelines analyzed detector data for *real* GW signals as well as simulated GW signals, called *injections*: GstLAL (Messick et al. 2017; Sachdev et al. 2019; Hanna et al. 2020; Cannon et al. 2020; Ewing et al. 2024; Tsukada et al. 2023; Sakon et al. 2024; Ray et al. 2023b; Joshi et al. 2025a,b), MBTA (Adams et al. 2016; Aubin et al. 2021; Andres et al. 2022; Alléné et al. 2025), PyCBC (Usman et al. 2016; Nitz et al. 2017, 2018; Dal Canton et al. 2021), and the cWB analysis (Klimenko et al. 2005, 2008, 2016; Tiwari et al. 2016; Drago et al. 2020; Klimenko 2022; Mishra et al. 2021, 2022, 2025). Injections were added to data at an artificially higher rate than observed GW signals, and were used to quantify the pipelines' sensitivities to GW signals (Essick et al. 2025; Abac et al. 2025b). The distribution of injections was chosen to enable efficient and accurate resampling to a wide range of astrophysically plausible populations (Essick et al. 2025). For all the analyses in this paper, the detection efficiency $\xi(\mathbf{\Lambda})$ is estimated using these injections through a Monte Carlo integral (Abac et al. 2025b; Tiwari 2018; Farr 2019).

3.4. *Source Properties*

Parameter estimation (PE) pipelines use Bayesian inference to estimate the properties of GW events (Abac et al. 2025b). The hierarchical Bayesian inference framework described in Section 2 requires as input PE samples from individual events. For events detected in O4a, we use samples drawn from the posterior distribution using the NRSur7dq4 (Varma et al. 2019) waveform approximant if available. If these are not available e.g., because the signal duration is too long to be analyzed by NRSur7dq4, we instead use a mixture of samples from the IMRPhenomXPHM_SpinTaylor (Pratten et al. 2021; Colleoni et al. 2025) and SEOBNRv5PHM (Ramos-Buades et al. 2023; Pompili et al. 2023) approximants. These are referred to as Mixed samples in the PE data products. More details about various choices made in the PE procedure can be found in Section 5 of Abac et al. (2025b) and Section 3 of Abac et al. (2025c).

For all events detected before O4a, we use the Mixed samples reported in the GWTC-3.0 (Abbott et al. 2023b) and GWTC-2.1 (Abbott et al. 2024) data releases. One exception is the BNS merger GW170817, for which we use samples obtained with the IMRPhenomPv2_NRTidal waveform approximant (Dietrich et al. 2019) and a prior allowing for large spin magnitudes (Abbott et al. 2019c). We also reweight these

samples to a distance prior that is uniform in comoving volume and source-frame time following the prescription in Appendix C of Abbott et al. (2021b).

While this work was in its final stages, a normalization error was discovered in the noise-weighted inner product used in the PE likelihood function (Abac et al. 2025b; Talbot et al. 2025b). While there is a version of the PE samples that account for the correct likelihood via a reweighting prescription (Abac et al. 2025b; Talbot et al. 2025b), we do not use these samples in this work. Further, for candidates detected during the first three observing runs, we discovered that incorrect priors were used when marginalizing over the uncertainty in the calibration of the LIGO detectors (Abac et al. 2025b). Preliminary re-analysis indicates that for each candidate the impact of this error in marginalization is small, and we expect the impact on our population analyses to be negligible compared to other sources of systematic error.

4. BINARY MERGER POPULATION ACROSS ALL MASSES

We begin our analysis of the astrophysical distribution of merging compact binaries with a joint analysis of all events discussed in Section 3—BNSs, NSBHs, and BBHs—without distinguishing between these different source classes. This allows for a broad look at the complete population, self-consistent measurements of the merger rates in each binary source class, and an analysis of the population at the transition between NSs and BHs. As mentioned in Section 3, we adopt a uniform detection threshold of FAR $< 0.25\ \mathrm{yr}^{-1}$ to ensure a high catalog purity of NS systems, where we have fewer detections and so are more sensitive to non-astrophysical false-alarm contaminants.

4.1. *The Mass Spectrum of Compact Binaries*

In Figure 1 we show the joint primary and secondary-mass distributions, inferred using a strongly modeled and a weakly modeled approach. Our FullPop-4.0 strongly modeled approach is modified from the Power law + Dip + Break analysis of the previous GWTC-3.0 catalog (Fishbach et al. 2020; Farah et al. 2022; Abbott et al. 2023a; Mali & Essick 2025); see Appendix B.1 for a model description. Our weakly modeled analysis approximates the m_1–m_2 space with a BGP (Mandel et al. 2017; Abbott et al. 2023a; Ray et al. 2023a, 2024); see Appendix C.2 for further discussion. The FullPop-4.0 model uses the default models in Table 1 for the redshift and component spins, and for NS masses ($m < 2.5\,M_\odot$) the spin magnitude distribution is truncated over the range $\chi \in [0, 0.4]$. The BGP analysis fixes the Power Law Redshift evolution to $\kappa = 3$ (see Appendix B.4), and the spin distribution to be uniform in magnitude (again truncated over $\chi \in [0, 0.4]$ for NS masses) and isotropic in orientation.

We observe an enhanced merger rate around $m_1 \sim m_2 \lesssim 2\,M_\odot$, representing BNS systems, an additional subpopulation at unequal masses $m_1 \sim 9\,M_\odot$ and $m_2 \lesssim 2\,M_\odot$ consistent with NSBHs, and a third BBH subpopulation at $m_1, m_2 \gtrsim 9\,M_\odot$. In the upper triangle of Figure 1, we show the fractional uncertainty as a function of mass, a unitless quantity $\Delta\mathcal{R}/\mathcal{R}$ defined as the 95th - 5th percentile uncertainty divided by the median merger rate. The rate is best constrained at equal masses, where the majority of mergers are observed, and in the BBH range 10–40 $M_\odot$. Due to fewer observations, the uncertainty is larger for BNS and NSBH systems.

In Figure 2, we show the marginalized primary and secondary-mass distributions for our strongly and weakly modeled reconstructions of the merger rate, focusing on the transition between NSs and BHs. The models are consistent within uncertainties, indicating that systematic error from model assumptions are smaller than the statistical uncertainties. The most notable difference between the FullPop-4.0 and BGP results is the large uncertainty in the BGP primary-mass distribution relative to the secondary mass. There are significantly more observed CBCs with light secondary masses $m_2 \lesssim 15\,M_\odot$ than primary masses $m_1 \lesssim 15\,M_\odot$ (simply by the definition $m_2 < m_1$), and so the data-driven BGP is more uncertain for small primary masses, while the FullPop-4.0 results are more model driven.

At $m_2 = 5\,M_\odot$, the FullPop-4.0 model adopts an alternative pairing function to more naturally distinguish the pairing behavior of BNS and NSBH systems from that of BBH systems, introducing a discontinuity in the secondary-mass distribution (Figure 2) at the transition point. We discuss the behavior of the population at the transition from NSs to BHs and the astrophysical insights below in Section 4.3.

4.2. *Merger Rates*

Since we self-consistently include all events in the catalog, our measurements of the binary merger rates are robust to events which straddle different source classes. We calculate the rates of BNS, NSBH, and BBH mergers by assuming any object with mass $1\,M_\odot < m < 2.5\,M_\odot$ is a NS, and any object with mass $m > 2.5\,M_\odot$ is a BH. In all models, we assume the rate evolves with redshift in a manner that is uncorrelated with mass (Fishbach

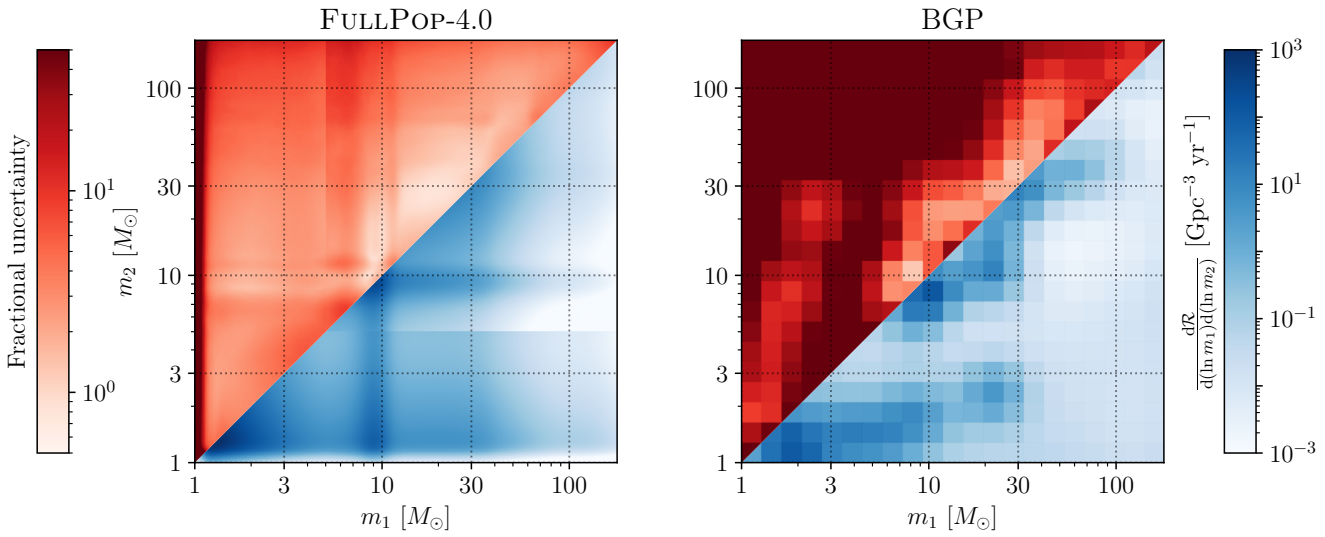

Figure 1. The complete mass spectrum, inferred using the strongly modeled FullPop-4.0 and weakly modeled BGP analyses. In the lower triangle of the m_1–m_2 plane (where $m_1 > m_2$ by definition), we show median merger rate density across primary and secondary masses. At $m_2 = 5\,M_\odot$, the FullPop-4.0 model transitions to an alternative pairing function to allow for NSBHs, hence the discontinuity at $m_2 = 5\,M_\odot$. In the upper triangle we show the fractional uncertainty in the merger rate reflected across the diagonal. The fractional uncertainty $\Delta\mathcal{R}/\mathcal{R}$ is difference of the 95th and 5th percentile values divided by the median merger rate. Both the FullPop-4.0 and BGP models show similar broad features: a population of BNSs at primary and secondary masses $\lesssim 2\,M_\odot$, a population of NSBHs at primary mass $\sim 9\,M_\odot$ and secondary mass $\lesssim 2\,M_\odot$, and finally the population of BBHs with primary and secondary masses $\gtrsim 9\,M_\odot$. The uncertainties are the smallest in the $\sim 9\,M_\odot$ and $\sim 30\,M_\odot$ BBH peaks.

Table 2. Merger rates in units $\mathrm{Gpc^{-3}\,yr^{-1}}$ in different mass ranges, according to the FullPop-4.0 and BGP models.

	BNS	NSBH	BBH	NS–Gap	BH–Gap	Full
	$m_1 \in [1, 2.5]\,M_\odot$	$m_1 > 2.5\,M_\odot$	$m_1 > 2.5\,M_\odot$	$m_1 \in [2.5, 5]\,M_\odot$	$m_1 > 2.5\,M_\odot$	$m_1 > 1\,M_\odot$
	$m_2 \in [1, 2.5]\,M_\odot$	$m_2 \in [1, 2.5]\,M_\odot$	$m_2 > 2.5\,M_\odot$	$m_2 \in [1, 2.5]\,M_\odot$	$m_2 \in [2.5, 5]\,M_\odot$	$m_2 > 1\,M_\odot$
FullPop-4.0	89^{+159}_{-67}	23^{+20}_{-13}	19^{+7}_{-5}	$6.7^{+9.2}_{-4.6}$	$2.5^{+3.4}_{-1.6}$	130^{+160}_{-80}
BGP	49^{+121}_{-42}	30^{+54}_{-21}	19^{+7}_{-5}	15^{+46}_{-14}	$1.5^{+4.9}_{-1.3}$	110^{+130}_{-60}
Merged	7.6–250	9.1–84	14–26	1.2–61	0.2–6.3	49–300
Simple Uniform BNS	13–170	–	–	–	–	–

Note—For BNS systems, we also estimate the rate assuming a Simple Uniform BNS model. We show rates of BNS, NSBH, and BBH assuming objects with mass $m \in [1, 2.5]\,M_\odot$ are NSs and $m > 2.5\,M_\odot$ are BHs. We also show rates within the purported lower-mass gap between astrophysical NSs and BHs, according to these models. In the third row, we show the merged estimates, taking the union of the 90% credible intervals for the FullPop-4.0 and BGP models, in order to account for model systematics. We quote merger rates at redshift $z = 0$.

et al. 2018; Abbott et al. 2023a). We quote merger rates in the local Universe, at redshift $z = 0$.

Our strongly-modeled FullPop-4.0 and weakly-modeled BGP approaches infer different rates over the m_1–m_2 space, and hence different rates in each binary source class. To marginalize over the systematic modeling uncertainty, we take the union of both 90% credible intervals. **The rates thus obtained are 7.6–250 $\mathrm{Gpc^{-3}\,yr^{-1}}$ for BNS mergers, 9.1–84 $\mathrm{Gpc^{-3}\,yr^{-1}}$ for NSBH mergers, and 14–26 $\mathrm{Gpc^{-3}\,yr^{-1}}$ for BBH mergers.** In Table 2, we show rates in these classes using different models and within purported mass gaps.

The BNS merger rate measurements may be sensitive to assumptions about NS pairing, and so we also estimate the BNS merger rate assuming a simple, fixed

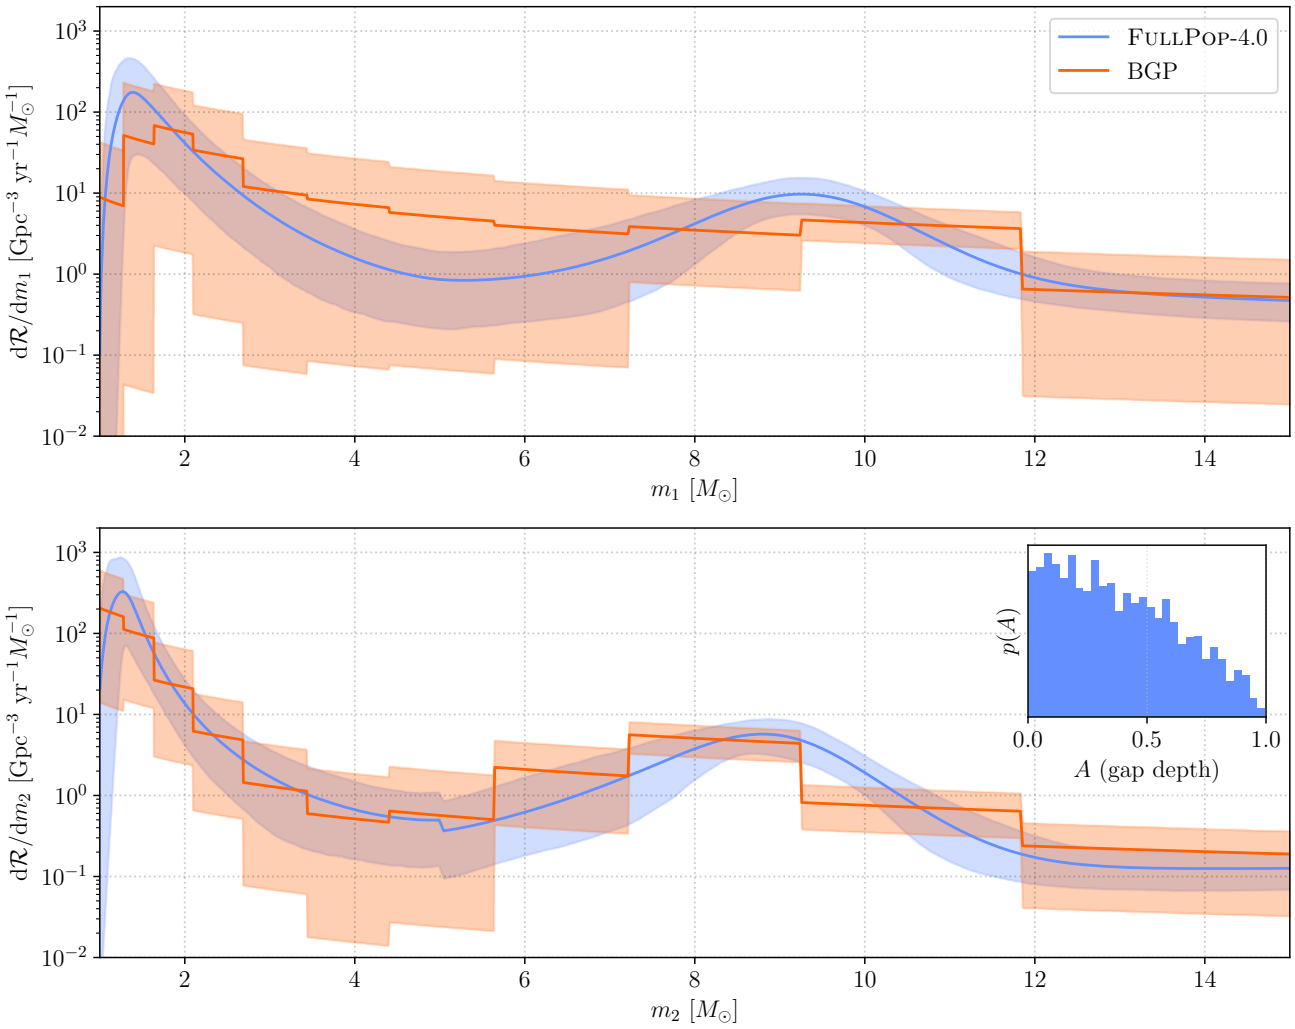

Figure 2. A comparison of the merger rate at redshift $z = 0$, as a function of component mass for the FullPop-4.0 and BGP models. In the upper panel, we show the merger rate as a function of primary mass, marginalized over secondary mass. In the lower panel, we show the merger rate as a function of the secondary mass, marginalized over the primary mass. At $m_2 = 5\,M_\odot$, the FullPop-4.0 model transitions to an alternative pairing function to allow for NSBHs, hence the discontinuity in the secondary mass. The median of the inferred merger rate is shown with a solid line, and the 90% credible interval is shown in the shaded region. The BGP weakly modeled approach has larger uncertainties due to increased flexibility. Both the strongly modeled and weakly modeled approaches find local maxima in the merger rate at $m_2 \sim 1 - 2\,M_\odot$ and $m_2 \sim 6.5 - 9\,M_\odot$ in the secondary-mass distribution. However, the BGP model does not confidently recover the low mass peaks in the primary-mass distribution. In the inset, we show the FullPop-4.0 inferred gap depth parameter A (see Appendix B.1), where $A = 1$ (a completely empty gap) is disfavored.

population. We assume a uniform mass distribution between $1\,M_\odot$ and $2.5\,M_\odot$ for the component masses, isotropically distributed spins with uniform spin magnitudes below 0.4, and a merger rate uniform in comoving volume up to $z = 0.15$. Under this fiducial model (denoted Simple Uniform BNS in Table 2), we infer a BNS merger rate of 13–170 $\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$.

The estimates in Table 2 are consistent with our previous analysis (Abbott et al. 2023a) and the uncertainties on the rate in each source class have generally decreased due to our larger catalog size. Although it is within uncertainties, our inferred merger rate for BNS systems has notably decreased by a factor of ~ 2 (cf. PDB (pair) and BGP in Table 2 of Abbott et al. 2023a). This is a result of the improved detector range and observing time together with the lack of new BNS detections.

4.3. *The Neutron Star–Black Hole Transition*

Electromagnetic (EM) observations have previously suggested the existence of a mass gap between NSs and

BHs (Bailyn et al. 1998; Ozel et al. 2010; Farr et al. 2011). On the lower end of the gap, nonrotating NS masses are bounded by a physical limit, the Tolman–Oppenheimer–Volkoff (TOV) mass (e.g., Kalogera & Baym 1996). Astrophysical observations, heavy-ion collision experiments, and modeling of the dense matter equation of state (EoS) at nuclear densities bound $M_{\rm max,TOV} \sim 2.2$–$2.5\,M_\odot$ (Landry et al. 2020; Dietrich et al. 2020; Legred et al. 2021; Huth et al. 2022; Ai et al. 2023; Dittmann et al. 2024; Rutherford et al. 2024; Koehn et al. 2025), and studies on the remnant in the BNS merger GW170817 (Abbott et al. 2017) place limits in the range $\lesssim 2.3\,M_\odot$ (Margalit & Metzger 2017; Rezzolla et al. 2018; Ruiz et al. 2018; Abbott et al. 2020b; Nathanail et al. 2021).

On the upper end of the gap, EM observations historically identified a dearth of BHs in the range ~ 3–$5\,M_\odot$ (Bailyn et al. 1998; Ozel et al. 2010; Farr et al. 2011), hinting at an astrophysical mass gap between NSs and BHs, or perhaps a selection effect obscuring such objects. More recently, observations of noninteracting binary systems (Thompson et al. 2018; Jayasinghe et al. 2021) and radio pulsar surveys (Barr et al. 2024) suggest the presence of a population of compact objects within the gap. Indeed, the GW events GW190814 (Abbott et al. 2020e) and GW230529 (Abac et al. 2024) are further evidence that the transition between NSs and BHs is populated, albeit sparsely.

With the additional GWTC-4.0 data, the GW picture of the purported lower-mass gap is becoming clearer and structures around the transition from NSs to BHs are emerging. In both our strongly modeled FullPop-4.0 and weakly modeled BGP analyses, we find evidence for **a prominent pair of peaks at NS masses $\sim 1.5 M_\odot$ and at BH masses $\sim 9\,M_\odot$ on each side of the lower-mass gap. A completely empty gap between NSs and BHs is disfavored.** We cannot rule out the existence of extremely narrow gaps in the compact object spectrum, though such a feature requires fine tuning of the supernova explosion mechanism, fallback, binary interactions or other physical processes (e.g., Fryer et al. 2012; Belczynski et al. 2012).

Our FullPop-4.0 model detects a peak in the BH merger rate at $9.02^{+0.41}_{-1.21}\,M_\odot$. We find that merging NSs represent the global maximum of the compact object mass spectrum at $1.23^{+0.14}_{-0.15}\,M_\odot$. At the transition from NSs to BHs, FullPop-4.0 allows for an additional suppression in the merger rate, parameterized by a gap depth parameter A where $A = 1$ corresponds to an absolute gap with zero mergers and $A = 0$ corresponds to no additional suppression. We show the measurement on A in the inset in Figure 2. A is consistent with zero, so the lower and upper bounds of the gap (e.g., the maximum NS mass and the minimum BH mass) are not measured away from the prior.

Earlier precursor analyses to FullPop-4.0 showed that the merger rate between ~ 3–$5\,M_\odot$ is likely suppressed relative to a power-law continuum (Fishbach et al. 2020; Abbott et al. 2021a; Farah et al. 2022), using older datasets. In Abbott et al. (2023a), we used a strongly modeled analysis on GWTC-3.0 (compare the PDB model to FullPop-4.0) to argue that the transition from NSs to BHs is likely suppressed, but also may be partially filled. FullPop-4.0 measures a similar mass spectrum to these previous analyses, but reintreprets the data as evidence for a rise at $\sim 9\,M_\odot$, and no *additional* suppression between the NS and BH peaks (Mali & Essick 2025 make similar conclusions on GWTC-3.0).

We corroborate our strongly modeled FullPop-4.0 conclusions with a weakly modeled BGP approach, and our constraints are improved relative to a similar BGP-based model in Abbott et al. (2023a). Unlike the FullPop-4.0 model, however, the BGP analysis does not distinguish between peaks, gaps, or a continuum: as a weakly modeled approach, the shape of the population is inferred directly without an associated interpretation.

The BGP model infers a nonzero merger rate between the NS to BH peaks consistent with an underlying continuum, and no evidence for a sharp gap. However, the smoothing kernel in the BGP model makes it *a priori* less sensitive to deep gaps, so we cannot rule them out either. The BGP model also observes NS and BH peaks. We quantify the significance as the fraction of hyperparameter samples where the merger rate density is higher than the merger rate in the neighboring bins. The merger rate maximizes for NSs in the range $1\,M_\odot \leq m_2 \leq 2\,M_\odot$ at 99% credibility and the low-mass BH peak occurs in the secondary mass $7.5\,M_\odot \leq m_2 \leq 9\,M_\odot$ bin at 97% credibility.

As our knowledge improves about the compact object population at the boundary between NSs and BHs, we stand to learn about supernova physics (e.g., Burrows & Vartanyan 2021) and the formation mechanisms of the heaviest NSs and lightest BHs (Abac et al. 2024, and references therein). If future catalogs continue to disfavor a completely empty gap, the standard picture of rapid core-collapse supernovae —which features a sharp transition in remnant masses; successful explosions leave NSs with $m \lesssim 2\,M_\odot$ and failed explosions promptly collapse to BHs with $m \gtrsim 5\,M_\odot$ (e.g., Fryer & Kalogera 2001; Fryer et al. 2012)— may require modifications to include fallback, slower instability growth, or stochasticity (e.g., Fryer & Kalogera 2001; Fryer et al. 2012;

Belczynski et al. 2012; Sukhbold et al. 2016; Ertl et al. 2019; Mandel & Müller 2020). Fallback of stellar material can produce black holes from the maximum neutron star mass to the lightest BH mass ($\sim 6\,M_\odot$) in the rapid implosion scenario (e.g., Ertl et al. 2019), or a slower instability growth timescale could allow the proto-NS to accrete enough mass before the explosion to populate the mass gap (e.g., Belczynski et al. 2012; Olejak et al. 2022; Fryer et al. 2022). Another possibility is that stochasticity in the stellar evolution and supernovae smooths out the remnant mass distribution and occupy the lower-mass gap (e.g., Mandel & Müller 2020; Mandel et al. 2020). Alternatively, the gap between NSs and BHs may be populated by a pollution mechanism, such as the remnants of BNS or white dwarf collisions which participate in further hierarchical mergers (e.g., Gupta et al. 2020; Ye et al. 2020, 2024; Barr et al. 2024; Mahapatra et al. 2025a) or other exotic scenarios like primordial black holes (e.g., Clesse & Garcia-Bellido 2022) or gravitationally lensed events, which could be mistaken for mass gap objects (e.g., Bianconi et al. 2023; Janquart et al. 2024; Farah et al. 2025a).

5. POPULATION PROPERTIES OF MERGERS CONTAINING NEUTRON STARS

In O4a, GW230529 (Abac et al. 2024) is the only NS-containing event identified with a FAR $< 0.25\,\mathrm{yr}^{-1}$. We do not include the NSBH candidate GW230518 identified during the ER preceding O4a. Additionally, no coincident EM counterparts were identified for triggers (i.e., preliminary candidate GW signals flagged by the detection pipelines when their ranking statistic exceeded the alert threshold) during O4a based on follow-up efforts conducted by EM telescopes. Thus, conclusions presented in previous papers (Abbott et al. 2023a; Abac et al. 2024) about the population properties of BNS and NSBH mergers remain largely unchanged.

We adopt models and methods consistent with those used in GWTC-3.0 (Abbott et al. 2023a). Specifically, we adopt two different models for the NS mass distribution, one in which masses are assumed to be Gaussian distributed, and another in which they are assumed to follow a power law (Abbott et al. 2023a). These are referred to as the Peak model and the Power model respectively. We assume that the redshift evolution of the merger rate is fixed, and that the spins are distributed following the PE prior (Abac et al. 2025b). A uniform prior is used for hyperparameters of the population, with the condition $m_{\rm min} \leq \mu \leq m_{\rm max}$ (where μ is the mean of the Gaussian bump in the Peak model), assuming $m_{\rm max}$ does not exceed $M_{\rm max,TOV}$ (the maximum permissible NS mass as expected from the TOV limit), as detailed in Appendix B.

When assumed to follow the Power model, NS masses favor a mass distribution with a power-law slope constrained to $\alpha = 7.7^{+5.1}_{-5.5}$. The inference from the Peak model is broad, with largely unconstrained peak width $\sigma = 0.68^{+1.2}_{-0.45}\,M_\odot$ and location $\mu = 1.4^{+0.48}_{-0.25}\,M_\odot$. While these results may hint at a peak emerging near $1.4\,M_\odot$, it is much broader than the relatively sharp peak in the mass distribution of Galactic NS systems (Farrow et al. 2019; El-Badry et al. 2024). Our inferred NS mass distribution remains broad, with greater support for high-mass NSs.

As no new NSBHs beyond GW230529 were confidently observed in O4a, the population-level results in Abac et al. (2024) produced using the NSBHPop model (Biscoveanu et al. 2022b) remain unchanged. Specifically, our inferred minimum BH mass in NSBH systems remains $3.4^{+1.0}_{-1.2}\,M_\odot$ (Abac et al. 2024).

6. BINARY BLACK HOLE POPULATION

In this section, we analyze the astrophysical population of BBHs using data from GW events with a FAR $< 1\,\mathrm{yr}^{-1}$. We only include events whose 1% lower limit on both component mass posteriors (under the PE priors) is larger than $3\,M_\odot$. There are 84 events from O4a that meet this criterion in addition to the 69 BBH events from GWTC-3.0, which brings the total number of BBH events passing our FAR threshold to 153. Of the events considered in O4a, only GW230529 does not meet this criterion and is not considered in the analyses below (see Section 3).

6.1. *Primary Mass*

In this section, we illustrate the main findings of the strongly and weakly modeled approaches using results from the Broken Power Law + 2 Peaks (see Appendix B.3) and B-Spline (see Appendix C.1) models, respectively. We chose the Broken Power Law + 2 Peaks model as our fiducial mass model because it performed the best in our model comparison study. It was first introduced by Callister & Farr (2024) to describe the GWTC-3.0 mass distribution. Results from a selection of other models considered in the model comparison study can be found in Appendix D.1, while models that employ a weakly modeled approach can be found in Appendix D.4, both of which include comparisons to the fiducial Broken Power Law + 2 Peaks and B-Spline models. The strongly modeled approach employs the default models listed in Table 1 for the other source parameters, namely the Power Law Redshift and Gaussian Component Spins models.

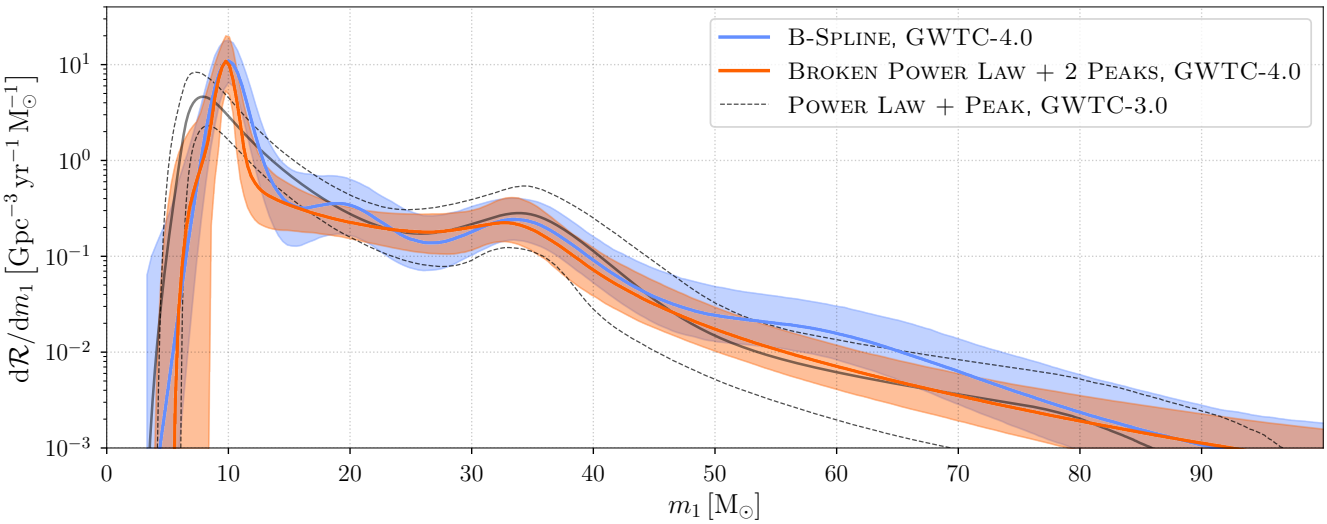

Figure 3. Differential merger rate as a function of primary mass (evaluated at $z = 0.2$) of the Broken Power Law + 2 Peaks model (orange) and B-Spline model (blue) compared to the Power Law + Peak model from GWTC-3.0 (Abbott et al. 2023a). The solid lines indicate the posterior medians and the shaded regions show the 90% credible interval of each model. Comparing these results, it is clear that a single power law is a poor description of the the low-mass end of the spectrum.

We identify a global peak at $\sim 10\,M_\odot$ and find that it is robust to model variations. A Gaussian-like peak at $\sim 10\,M_\odot$ was first identified by Tiwari & Fairhurst (2021) and Edelman et al. (2022) using GWTC-2.0 and later by the weakly modeled approaches in Abbott et al. (2023a) using GWTC-3.0. Figure 3 shows the rate $d\mathcal{R}/dm_1$ as a function of primary mass at $z = 0.2$ inferred by these models compared to the Power Law + Peak (see Appendix B of Abbott et al. 2023a for a description of this model) inference with GWTC-3.0 from Abbott et al. (2023a).[1] The low-mass Gaussian component of the Broken Power Law + 2 Peaks model infers a global peak in the primary mass spectrum at $m_1 = 9.8^{+0.3}_{-0.6}\,M_\odot$ relative to the underlying broken power law continuum. The B-Spline model also recovers a global peak at $m_1 = 10.1^{+0.7}_{-0.7}\,M_\odot$. Both values are consistent with GWTC-3.0, which inferred a global peak at $m_1 = 10.2^{+0.3}_{-0.6}\,M_\odot$ with various weakly modeled approaches.

A broken power law is necessary to describe the continuum structure above $\sim 15\,M_\odot$ (see Figure 3). Below $\sim 35\,M_\odot$, the continuum of BH masses is well described by a power law with spectral index $\alpha_1 = 1.7^{+1.2}_{-1.8}$. Above $\sim 35\,M_\odot$, the continuum steepens, with a spectral index of $\alpha_2 = 4.5^{+1.6}_{-1.3}$ that is consistent with the Power Law + Peak result from GWTC-3.0 ($\alpha = 3.5^{+0.6}_{-0.6}\,M_\odot$). We find that $\alpha_2 > \alpha_1$ at 97.7% credibility. Though this continuum structure was identified by Callister & Farr (2024), it was not identified by the strongly modeled approaches in Abbott et al. (2023a), in part due to the limited flexibility of models employed in Abbott et al. (2023a). To illustrate this, in Figure 4 we present a re-analysis of GWTC-3.0 using the Broken Power Law + 2 Peaks model alongside the GWTC-4.0 result, which shows improved constraints across the full mass spectrum.

We identify a feature at $\sim 35\,M_\odot$. Using the strongly modeled approach, we find that this feature is consistent with either: (i) an over-density that peaks at $m_1 = 32.7^{+2.7}_{-6.5}\,M_\odot$ relative to an underlying broken power law (i.e., the Broken Power Law + 2 Peaks result), or (ii) a broken power law with break mass $m_{\rm break} = 34.1^{+3.8}_{-3.3}\,M_\odot$ (i.e., a broken power law without a second peak near the break). See Appendix D.1 for a more detailed discussion of this result. The former conclusion is supported by the B-Spline model, which exhibits a local peak at $m_1 = 33.5^{+2.2}_{-4.5}\,M_\odot$, consistent with other models that employ a weakly modeled approach in Appendix D.4.

Additional structure may be present in the mass spectrum. A bump near $\sim 20\,M_\odot$ is present in some of the weakly modeled approaches (see Figure 20 in Appendix D.4). We cannot conclude with the strongly modeled approach whether adding a third Gaussian component to the Broken Power Law + 2 Peaks model in this region is required by the data or not (as quantified by the log Bayes factor $\log_{10}\mathcal{B} = -0.34$ between the default model and one including a third

[1] We quote BBH merger rates at $z = 0.2$ to be consistent with Abbott et al. (2023a) and because $z \sim 0.2$ is where we best constrain the merger rate.

peak). This feature was first reported in an analysis of GWTC-2.0 (Tiwari & Fairhurst 2021) and was present in several analyses of GWTC-3.0 (Abbott et al. 2023a; Edelman et al. 2023; Tiwari 2023; Godfrey et al. 2023). Additionally, the B-Spline and other weakly modeled approaches show a rise in the merger rate relative to the Broken Power Law + 2 Peaks result in the $\sim 60\,M_\odot$ region, which can be seen in Figure 3 and Figure 20. A previous study of GWTC-3.0 found evidence for a similar feature in this region (Magaña Hernandez & Palmese 2025).

We do not place informative constraints on the individual parameters that govern low-mass smoothing for the Broken Power Law + 2 Peaks model. We caution against astrophysically interpreting the BBH primary mass distribution below $\sim 8\,M_\odot$ because of the bias that may be introduced by removing the probable NS-containing events in the manner described in Section 3. Removing such events is responsible for the discrepancy below $\sim 8\,M_\odot$ between Figure 2 and Figure 3.

GWTC-4.0 includes an exceptional high-mass BBH, GW231123 (Abac et al. 2025e), whose inferred component masses lie at the extreme upper end of those in our dataset. To check if GW231123 is an outlier with respect to the mass distribution of BBHs, we construct mock catalogs containing 153 detected events following the inferred population without GW231123, and calculate the distribution of maximum detectable BH masses. The total mass of GW231123 lies at the 87^{+6}_{-10}th percentile of this distribution. This shows that while GW231123 lies in the tail of the distribution, its total mass is consistent with the inferred mass spectrum; the degree of consistency is more than was the case with GWTC-3.0 data alone (Abac et al. 2025e).

The main compact binary formation scenarios (Mandel & Farmer 2022 and references therein)—isolated binary evolution and dynamical assembly in dense stellar environments—have both been shown to produce populations consistent with current observations (Mandel & Broekgaarden 2022). A peak near $m_1 \sim 10\,M_\odot$ is often predicted by isolated binary evolution models (Dominik et al. 2015; Belczynski et al. 2020; Giacobbo & Mapelli 2018; Wiktorowicz et al. 2019; Neijssel et al. 2019), while mass distributions from dynamical formation, such as in young and globular clusters, typically peak above $10\,M_\odot$ (Rodriguez et al. 2016a; Hong et al. 2018; Rodriguez et al. 2019; Banerjee 2021; Antonini & Gieles 2020). However, predictions from isolated and dynamical formation often overlap and can vary significantly based on the assumptions and methodologies used, making it difficult to conclude the origin of the observed catalog or constrain formation physics based on features in the marginal mass distributions.

A feature that may provide distinguishing power between the isolated and dynamical channels is the existence of an upper mass gap, in the range $45\,M_\odot \lesssim m \lesssim 120\,M_\odot$. Such a dearth would be consistent with the theorized *pair-instability mass gap*, arising from the complete disruption of massive stars due to runaway electron–positron pair production (Woosley 2017; Mapelli et al. 2020; Farmer et al. 2019). While the precise locations of the lower and upper edges of the pair-instability mass gap are sensitive to physical assumptions (Renzo et al. 2020; van Son et al. 2020; Woosley & Heger 2021; Shen et al. 2023; Winch et al. 2024), the feature tends to be a robust prediction of most stellar evolution models (Marchant et al. 2018; Marchant & Moriya 2020; Renzo et al. 2020; Woosley & Heger 2021). The gap may not be completely empty due to overmassive stellar envelope fallback or stellar mergers (Di Carlo et al. 2019, 2020; Mapelli et al. 2020; Kremer et al. 2020), hierarchical BBH mergers in clusters or in AGN environments (Mckernan et al. 2018; Rodriguez et al. 2019; McKernan et al. 2020; Yang et al. 2019; Mapelli et al. 2021; Antonini et al. 2019; Fragione & Silk 2020; Liu & Lai 2021; Martinez et al. 2020; Arca Sedda 2020; Mahapatra et al. 2021, 2025b, 2024), or even due to primordial BHs (Postnov & Yungelson 2014; Bird et al. 2016; Clesse & Garcia-Bellido 2022). The decrease in the merger rate above $m_1 \sim 35\,M_\odot$ seen in all models and the subsequent rise near $m_1 \sim 60\,M_\odot$ seen in the weakly modeled approaches may hint toward a polluted mass gap.

6.2. *Mass-Ratio*

Figure 5 shows the differential merger rate $\mathrm{d}\mathcal{R}/\mathrm{d}q$ evaluated at $z = 0.2$ inferred by the Broken Power Law + 2 Peaks and B-Spline models compared to the Power Law + Peak result from GWTC-3.0.[2]

We find that the mass-ratio distribution can be described by a power-law with an index $\boldsymbol{\beta_q = 1.2^{+1.2}_{-1.0}}$ that is consistent with GWTC-3.0 ($\beta_q = 1.1^{+1.7}_{-1.3}$) with a reduction in uncertainty. The B-Spline model shows a reduced rate above $q \gtrsim 0.8$ relative to the Broken Power Law + 2 Peaks result. More flexible parametrizations in the strongly modeled approach were explored (see also the other weakly modeled approaches in Figure 20, which do peak near unity), but model com-

[2] Models that utilize B-Splines infer the separable distribution $p(q)$ but the results of such models shown in this section are actually the conditional marginal distribution $p(q|m_2{>}3M_\odot)$. See Appendix C.1 for further details.

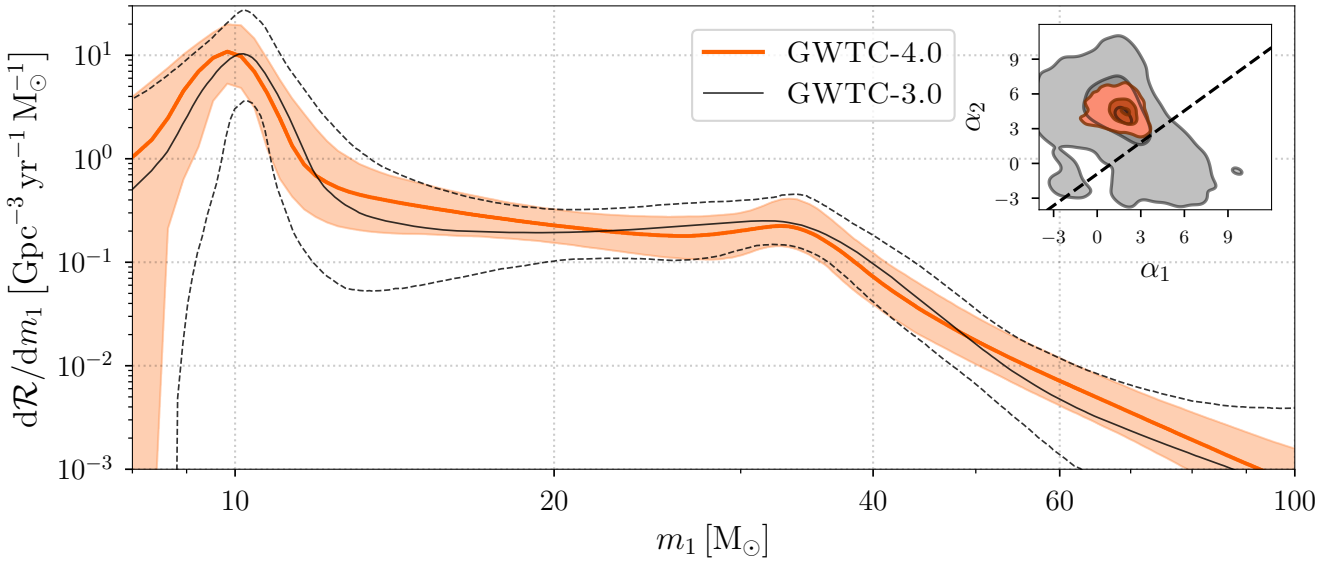

Figure 4. Differential merger rate as a function of primary mass (evaluated at $z = 0.2$) of the Broken Power Law + 2 Peaks model inferred with GWTC-4.0 compared to the same model applied to only GWTC-3.0 BBH events. The orange shaded region shows the 90% credible interval for GWTC-4.0 and the solid orange curve shows the posterior median, while the black dashed curves bound the 90% credible region for GWTC-3.0 and the solid black curve shows the posterior median. The inferred distribution is similar between catalogs, which highlights that the low-mass structure identified in GWTC-4.0 was present in GWTC-3.0. The inset figure shows the joint posterior of the broken power law index parameters α_1 and α_2 for GWTC-3.0 (gray) and GWTC-4.0 (orange), with the contours showing the $5^{\mathrm{th}}, 50^{\mathrm{th}}$, and 95^{th} percentiles. The black dashed curve in the inset indicates where $\alpha_1 = \alpha_2$.

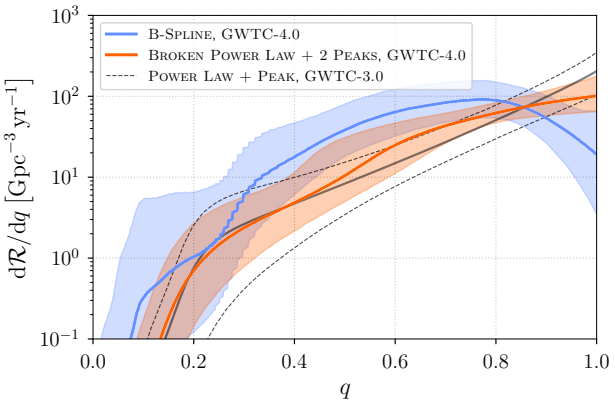

Figure 5. Differential merger rate as a function of mass-ratio (evaluated at $z = 0.2$) of the Broken Power Law + 2 Peaks model (orange) and B-Spline model (blue). The fiducial Power Law + Peak model from Abbott et al. (2023a) is included for comparison. Solid curves indicate posterior medians and the shaded (dashed) regions show 90% credible intervals.

parisons were inconclusive. Similarly, in GWTC-3.0, Godfrey et al. (2023) inferred a mass-ratio distribution peaked away from unity using a weakly modeled approach, while Rinaldi et al. (2025) found equal evidence for two strongly modeled approaches with very distinct behavior above $q \gtrsim 0.7$.

Models that incorporate correlations between source parameters have a greater potential to distinguish between formation channels than uncorrelated ones. We next present results from two different models that correlate features of the primary mass spectrum with different mass-ratio distributions.

BHs with masses $\sim$35 $M_\odot$ preferentially merge with other BHs of more equal mass relative to those in the underlying mass continuum. The Extended Broken Power Law + 2 Peaks model modifies the Broken Power Law + 2 Peaks model by allowing each primary mass mixture component (i.e., the broken power law and two Gaussian components) to be associated with a different power-law-mass-ratio distribution. Figure 6 shows the mass-ratio distribution for each mixture component. Specifically, each component infers a different power-law index β_q, with $\beta_q^{\mathrm{BP}} = 0.1^{+1.7}_{-1.3}$ for the broken power law, $\beta_q^{\mathrm{peak1}} = 1.6^{+9.6}_{-8.1}$ for the Gaussian component that captures the $\sim 10\,M_\odot$ peak, and $\beta_q^{\mathrm{peak2}} = 7.4^{+4.8}_{-5.1}$ for the second Gaussian component that captures the $\sim 35\,M_\odot$ feature. Critically, $\beta_q^{\mathrm{peak2}} > \beta_q^{\mathrm{BP}}$ at 97.3% credibility. Other studies of GWTC-3.0 have drawn similar conclusions about the pairing preferences of high mass BHs (e.g., Li et al.

2022; Baibhav et al. 2023; Sadiq et al. 2024; Galaudage & Lamberts 2025; Roy et al. 2025).

BHs with masses $\sim 10\,M_\odot$ may prefentially merge with lighter BHs. The region bounded by the solid blue curves in the bottom panel of Figure 6 shows the mass-ratio distribution of the $\sim 10\,M_\odot$ peak inferred with the ISOLATED PEAK model. This model is a mixture of a Gaussian peak and a B-Spline (continuum) in primary mass, and the mass ratio and spin distributions are inferred separately for each mixture component with B-Splines (see Appendix C for further details). The Gaussain peak is inferred at $\sim 10\,M_\odot$ and its mass-ratio distribution exhibits a peak at $q = 0.74^{+0.13}_{-0.13}$, a feature that a single power law is unable to reproduce. This could explain the large uncertainty in $\beta_q^{\rm peak1}$ from the EXTENDED BROKEN POWER LAW + 2 PEAKS model. The solid blue curve in Figure 6 shows the mass-ratio distribution of masses outside of the $\sim 10\,M_\odot$ peak inferred by the B-Spline mixture component. Unlike the full population mass-ratio distribution inferred by the B-SPLINE model in Figure 5, this result does not possess a peak near $q \sim 0.8$, indicating that the peak seen in the full population is due largely to the events around $\sim 10\,M_\odot$. This mass feature was identified in GWTC-3.0 by Godfrey et al. (2023).

Most formation channels generally favor equal mass systems. For example, dynamical formation can produce systems with a wide range of mass ratios, but predicted distributions typically peak at unity (Rodriguez et al. 2016a; Torniamenti et al. 2024). Certain hierarchical mergers may not necessarily follow this trend, in particular mergers between first generation and second generation BHs have been shown to produce a mass-ratio distribution peaked near $q \sim 0.5$ (Rodriguez et al. 2019). Mass transfer during the contact phase of binaries formed via chemically homogeneous evolution (de Mink & Mandel 2016; Marchant et al. 2016) leads to a strong preference for equal mass-ratio systems, but this mechanism is thought to be important for binaries with $m_1 \gtrsim 10\,M_\odot$ (du Buisson et al. 2020; Zevin et al. 2021; Riley et al. 2021). Mass-ratio reversal within the stable mass transfer channel can lead to a peak in the mass-ratio distribution between $q \sim$ 0.6-0.8 (Neijssel et al. 2019; van Son et al. 2022b), which is qualitatively consistent with the mass-ratio distribution of the $\sim 10\,M_\odot$ inferred with the ISOLATED PEAK model. Stable mass transfer also predicts a peak near $\sim 10\,M_\odot$ that is robust to uncertainties in the metallicity-dependent star formation history (van Son et al. 2022a) and other physical uncertainties of the channel (van Son et al. 2022c).

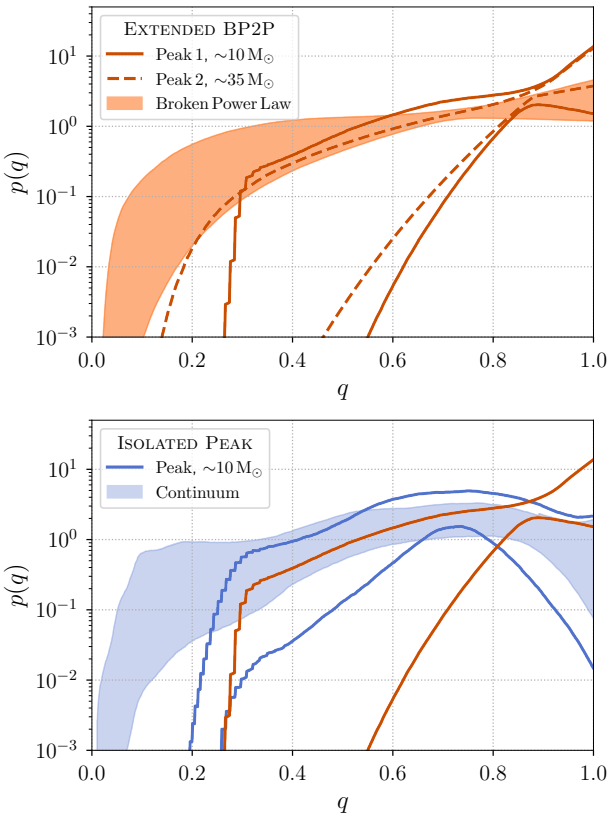

Figure 6. Top Panel: The 90% credible regions for the mass-ratio distribution of the $\sim 10\,M_\odot$ peak (solid orange curves), $\sim 35\,M_\odot$ peak (dashed curves), and continuum (orange shaded region) components of the EXTENDED BROKEN POWER LAW + 2 PEAKS model. Bottom Panel: The 90% credible regions for mass-ratio distribution of the $\sim 10\,M_\odot$ peak (solid blue curves) and the rest of the mass spectrum (blue shaded region) inferred with the ISOLATED PEAK model. The $\sim 10\,M_\odot$ peak mass-ratio distribution from the EXTENDED BROKEN POWER LAW + 2 PEAKS is included for comparison.

6.3. *Spin*

We next present the spin distribution of the BBH population through O4a. We model spins using two different parameterizations: the magnitudes χ_i and tilt angles θ_i (Section 6.3.1), and the effective inspiral spin $\chi_{\rm eff}$ and effective precessing spin $\chi_{\rm p}$ (Section 6.3.2). The effective inspiral spin $\chi_{\rm eff}$ is the mass-weighted average of the component spins *aligned* with the binary's orbital angular momentum (Racine 2008; Ajith et al. 2011). The effective precessing spin $\chi_{\rm p}$ characterizes the degree of relativistic precession caused by spin–orbit misalignment, capturing the effect of the *in-plane* spin components (Schmidt et al. 2011, 2012, 2015; Gerosa et al. 2021); these are defined in Equations 15 and 16 of Abac

et al. (2025a). More details about spin parameterization are given in Appendix B.5.

Previous analyses of the BBHs observed through O3 found that BH spin vectors tend to be small in magnitude ($\chi \lesssim 0.3$) and preferentially—but not exclusively—lie above the orbital plane ($\cos\theta > 0$; Abbott et al. 2019a, 2021a, 2023a). The new observations in GWTC-4.0 further support these conclusions and additionally suggest more structure in the component and effective spin distributions, enabled by changes to the models used in previous population analyses: (i) **the spin magnitude distribution has support at $\chi \approx 0$**, (ii) **the spin tilt distribution may peak away from perfect alignment with the orbital angular momentum**, and (iii) **the $\chi_{\rm eff}$ distribution is asymmetric about its peak**. We elaborate upon these new features in the following subsections.

6.3.1. *Spin Magnitudes and Tilts*

Spin magnitudes and tilt angles provide insight about BBH formation and evolution (e.g., Mandel & Farmer 2022; Mapelli 2020); we begin with spin magnitudes. If angular momentum transport in stars is efficient, stellar cores rotate slowly, resulting in small spin magnitudes for isolated BHs (Fuller & Ma 2019; Ma & Fuller 2019; Fuller et al. 2019). However, binary interactions can significantly influence BH spins through mechanisms such as tides (Hut 1981; Packet 1981; Zaldarriaga et al. 2018; Qin et al. 2018; Bavera et al. 2020; Mandel & Fragos 2020) and accretion (Hut 1981; Packet 1981; van den Heuvel et al. 2017; Neijssel et al. 2019; Steinle & Kesden 2021; Stegmann & Antonini 2021). If a BBH forms from a pre-existing isolated stellar binary, tidal interactions can spin up the progenitor of the second-born BH (Bavera et al. 2020; Qin et al. 2019), yielding spin magnitudes of $\chi \approx 0.2{-}0.4$ (Ma & Fuller 2023), although these can be reduced by stellar winds (Tout & Pringle 1992). Chemically homogeneous evolution—involving tidally locked, high-mass, low-metallicity, close binaries—may produce even larger spins (Mandel & de Mink 2016; de Mink & Mandel 2016). Alternatively, large spin magnitudes may point toward a hierarchical merger origin, where one or both BHs are remnants of previous BBH mergers (Rodriguez et al. 2019; Zhang et al. 2023; Doctor et al. 2019; Kimball et al. 2020; Payne et al. 2024; Gerosa & Fishbach 2021; Fishbach et al. 2017; Gerosa & Berti 2017; Mould et al. 2022b; Mahapatra et al. 2021, 2025b, 2024), which are predicted to have $\chi \sim 0.7$ (Lousto et al. 2010).

With this astrophysical context, we present our measurement of the spin magnitude distribution through O4a. The left column of Figure 7 shows the marginal distributions of χ using the strongly modeled approach (Gaussian Component Spins; blue) and weakly modeled approach (B-Spline; green). The Gaussian Component Spins model (Equation B26) describes the χ population as a truncated Gaussian distribution. This is a departure from the Default spin model of GWTC-3.0. There, a non-singular Beta distribution was used (Abbott et al. 2023a), which forces $p(\chi) = 0$ at $\chi = 0, 1$ and thus cannot measure contributions to the population at near-minimal or near-maximal spins. Allowing for more model flexibility at the χ distribution's boundaries is crucial, as there exists an ongoing discussion in the literature about whether or not there is an over-density of BBHs with $\chi \lesssim 0.01$ (Kimball et al. 2020; Galaudage et al. 2021; Callister et al. 2022; Tong et al. 2022; Mould et al. 2022a; Hussain et al. 2024). We find that the Gaussian Component Spins model is preferred over the Default spin model of GWTC-3.0 by $\log_{10}\mathcal{B} = 0.66$; additional parametric spin magnitude (and tilt) models are discussed in Appendix D.2. The widening of the 90% credible regions for the χ and $\cos\theta$ distributions at their boundaries under the B-Spline model seen in Figure 7 is a prior-driven effect common in spline modeling (e.g., Golomb & Talbot 2023; Edelman et al. 2023).

In GWTC-4.0, we constrain $p(\chi \approx 0) > 0$ under both the strongly and weakly modeled approaches. At 90% credibility, our recovered spin magnitude distribution peaks between $\chi = 0.01$–0.23, as measured by the μ_χ location parameter. Broadly, BH spins are inferred to be predominantly non-extremal, with the Gaussian Component Spins model finding that 90% of BHs having $\chi < 0.57$. The comparative dearth of observed BBHs with large spins disfavors a population dominated by second-generation BHs. However, the precise fraction of systems with large spin magnitudes is model dependent: the Gaussian Component Spins and B-Spline models only disagree at 90% credibility for $\chi > 0.83$. At $\chi = 0.8$, 0.9, and 1.0, their $p(\chi)$ distributions differ at the 85%, 96%, and 98% levels, respectively. While the Gaussian Component Spins model approaches $p(\chi) \sim 0$ at $\chi = 1$, the B-Spline model infers a larger, nearly flat distribution from $\chi = 0.6$ to 1. Similar behavior was found in GWTC-3.0 (Godfrey et al. 2023), and attributed to a subpopulation with near-uniform spin magnitudes. The discrepancy between our two models may arise from limited flexibility in the strongly modeled approach; a single truncated Gaussian cannot increase the probability at $\chi \gtrsim 0.8$ without also increasing it at $\chi \sim 0.2{-}0.5$. In GWTC-4.0, a handful of events do have preferentially large spin magnitudes, including GW231123 which has

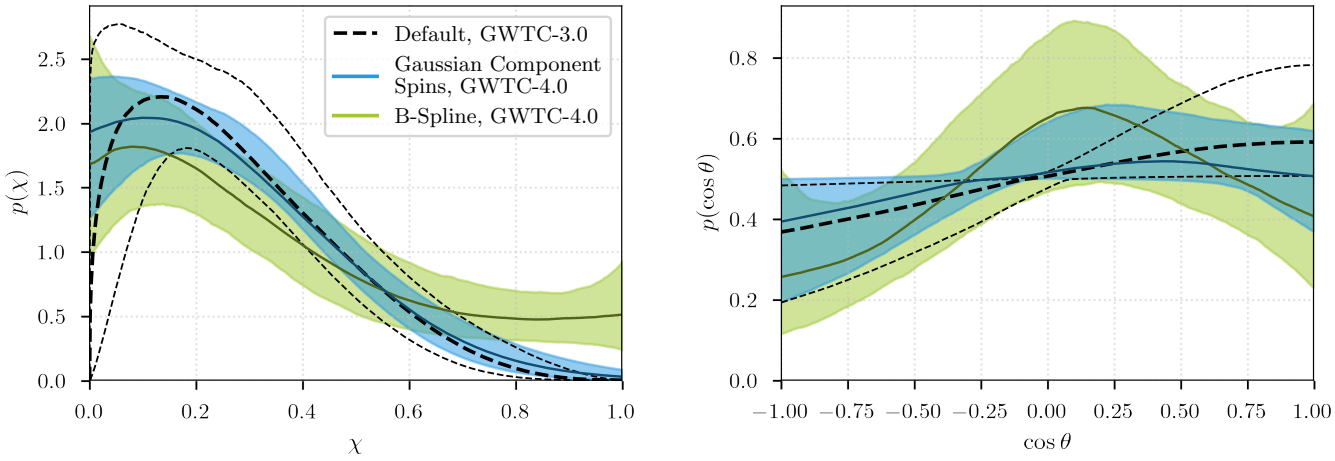

Figure 7. Marginal spin magnitude χ (left) and cosine tilt angle $\cos\theta$ (right) distributions under the Gaussian Component Spins (blue) and B-Spline (green) models. The results from the GWTC-3.0 Default are shown for comparison (black dashed); note that this model is different from what is used for GWTC-4.0. The GWTC-3.0 analysis employed a Beta distribution for spin magnitudes rather than a truncated normal, and fixed the location of the Gaussian component of the spin tilt angle distribution to $\cos\theta = 1$ rather than letting it vary freely in inference. The solid lines show the median of the inferred distributions from GWTC-4.0, and the shaded regions show their 90% credible regions. The thick dashed lines show the median from GWTC-3.0, while the thin dashed lines show its 90% credible region.

primary spin $\chi_1 > 0.5$ with high confidence (Abac et al. 2025e).

Under both the Gaussian Component Spins and B-Spline models, the results presented in Figure 7 assume that the spin magnitudes are independently and identically distributed (IID), meaning that the function describing the population distribution is factorizable in terms of χ_1 and χ_2, and the two are described by the same set of hyperparameters. The data prefer identically distributed spin magnitudes over those which are non-identically distributed, cf. Table 12. However, mathematically, the primary and secondary spins cannot be IID if the purported correlation between q and $\chi_{\rm eff}$ is true (Farr & Farr 2025). We probe this correlation in Section 6.5.2, and find support for its existence, albeit with evidence that has diminished since GWTC-3.0. Thus, we interpret the Bayes factor in favor of IID spins as a statement that, under the Gaussian Component Spins model, we cannot yet say with statistical certainty that the spins are not identically distributed. This statement is likely driven by the fact that individual-event spin magnitude posterior distributions are typically wide, making χ_1 and χ_2 hard to distinguish. In general, χ_1 and χ_2 are not expected to be identically distributed in nature; if q and $\chi_{\rm eff}$ are indeed correlated, more informative χ_i measurements may eventually reveal their non-identical nature. Figure 17 in Appendix D.2 presents posteriors on the Gaussian Component Spins hyperparameters under the assumption that spin magnitudes (and tilts) are identically versus non-identically distributed.

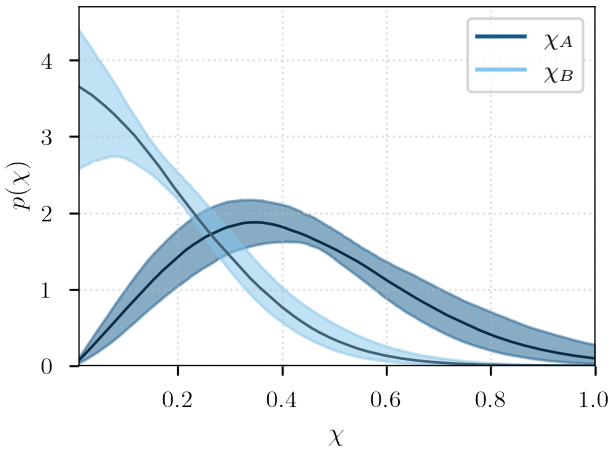

Figure 8. Larger (χ_A, lighter blue) and smaller (χ_B, darker blue) spin magnitude distributions, derived from imposing order statistics on the Gaussian Component Spins χ distribution shown in Figure 7. The solid lines show the median of each inferred distribution, and the shaded regions show the 90% credible intervals.

Next, Figure 8 shows the inferred spin magnitude distributions if, rather than sorting by the more versus less massive BH, we instead sort by the BH with the larger (subscript A) and smaller (subscript B) spin magnitude (Biscoveanu et al. 2021). The magnitudes χ_A and χ_B are derived quantities and are not fit independently: to generate their distributions, we take

results from the Gaussian Component Spins model and impose order statistics, assuming that χ_A is the larger of two draws from the $p(\chi)$ distribution in the left panel of Figure 7 and χ_B is the smaller (Abbott et al. 2023a; Biscoveanu et al. 2021). This analysis does not assume any sort of pairing function between spins. Spin sorting offers an alternative way to visualize and interpret the spin information from the IID model. The more rapidly spinning component has a wide spin magnitude distribution, with a population predictive distribution (PPD) peaking at $\chi_A = 0.35$, and support up to $\chi_{A,99\%} = 0.92^{+0.04}_{-0.08}$ (value of χ_A at which each $p(\chi_A)$ trace reaches its 99th percentile, serving as a proxy for the distribution's maximum). The vanishing probability at $\chi_A = 0$ is a Jacobian effect of the order statistics. However, if both BHs had $\chi \approx 0$, the χ_A distribution would be much more strongly peaked at small values (Szemraj & Biscoveanu 2025), meaning that spin-sorting results on GWTC-4.0 indicate that **at least one BH per binary has $\chi \gtrsim 0$**. The more slowly spinning component, on the other hand, is consistent with a narrower distribution peaking at $\chi_B = 0$, and only has support up to $\chi_{B,99\%} = 0.6^{+0.09}_{-0.07}$. The observation that spin χ_B magnitudes are preferentially small could indicate small natal spins for at least one of the two BHs, while the fact that the population is consistent with only one BH per binary having large spin supports the existence of some variety of spin-up mechanism in BBH evolution.

We next turn to the spin tilt distribution. Many authors argue that if BBHs are formed primarily in the isolated binary scenario, large tilt angles are hard to explain without invoking large supernovae kicks and inefficient tides (Kalogera 2000; Gerosa et al. 2018; Steinle & Kesden 2021; Wysocki et al. 2018; Stevenson 2022; Callister et al. 2021a), although others claim that, depending on the specifics of poorly understood supernovae physics, even small kicks can misalign a binary (Baibhav & Kalogera 2024; Tauris 2022). Complete anti-alignment ($\cos\theta = -1$) is a possible result of mass transfer in the isolated channel (Stegmann & Antonini 2021). On the other hand, BBHs formed dynamically in stellar clusters are predicted to have istropically distributed spin orientations, as there is no *a priori* preferential spin direction in these environments (Rodriguez et al. 2015, 2016a, 2018; Farr et al. 2017). However, recent studies indicate that mechanisms could possibly exist for dynamically formed BBHs to have slight preference for $\cos\theta > 0$ (Trani et al. 2021; Banerjee et al. 2023; Kıroğlu et al. 2025), especially for BBHs formed dynamically in the disks of active galactic nuclei (Wang et al. 2021; McKernan et al. 2022).

The right-hand column of Figure 7 shows the marginal distribution of the cosine of the tilt angle, $\cos\theta$. The $\cos\theta$ population under the Gaussian Component Spins model is a mixture between isotropic and truncated Gaussian sub-populations (Equation B27). Following GWTC-3.0 and earlier population analyses, we assume that spin tilt angles are nonindependently but identically distributed (NID) under the Gaussian Component Spins model, meaning that while $\cos\theta_1$ and $\cos\theta_2$ are described by the same hyperparameters, the population distribution is *not* separable in terms of the two: we require that both BHs in a binary are drawn from the same sub-population (either the isotropic or the truncated Gaussian). NID tilts are favored by the data over non-identical distribution, cf. Table 12. The B-Spline model naturally assumes spin tilt angles are IID, as it does not probe separate tilt sub-populations.

The Default model of GWTC-3.0 fixed the location of the Gaussian sub-population at exact spin–orbit alignment (Vitale et al. 2017; Talbot & Thrane 2017; Abbott et al. 2023a), making it a half-Gaussian peaking at $\cos\theta = 1$. In GWTC-4.0, the strongly and weakly modeled approaches both find that **the spin tilt distribution may peak away from exact spin–orbit alignment**. The Gaussian Component Spins model finds that the $\cos\theta$ distribution reaches its maximum between -0.36 and 0.94 at 90% credibility; consistently, the B-Spline model finds the peak to lie between -0.23 and 0.96. The possibility that the tilt angle distribution peaks away from alignment is also found in GWTC-3.0 under various models which permit such a feature (Vitale et al. 2022; Edelman et al. 2023; Golomb & Talbot 2023). It is not impossible, however, for the peak of a $\cos\theta$ distribution to be inferred away from alignment even when the true underlying population does peak at $\cos\theta = 1$ (Vitale & Mould 2025). The inferred peak of the GWTC-4.0 $\cos\theta$ distribution is more strongly constrained away from unity than those spuriously found with simulated catalogs of the same size (c.f., Vitale & Mould 2025, Figure 4).

Under both the strongly and weakly modeled approaches, the $\cos\theta$ distribution shown in Figure 7 has support across a wide range of spin tilts, with a slightly larger fraction of positive $\cos\theta$ compared to negative. Given this broad support, we next ask the question: what is the lowest spin tilt angle that is absolutely required to fit GWTC-4.0 reasonably? To probe this, we use a model which is similar to the Gaussian Component Spins model but with $p(\cos\theta < t_{\rm min}) = 0$ for an inferred value $t_{\rm min}$, where $t \equiv \cos\theta$ (Tong et al. 2022; Callister et al. 2022); see Equation (B28). We find $t_{\rm min} = -0.71^{+0.15}_{-0.19}$ at 90% credibility. That the

posterior for the minimum required $\cos\theta$ is inconsistent with -1 but remains confidently negative is a somewhat unexpected astrophysical result. If most BBHs form in isolation, a minimum tilt cutoff is plausible but would be expected to be closer to zero, potentially even positive. Conversely, dynamically formed BBHs should exhibit isotropic tilts extending down to $\cos\theta = -1$. Our observed tilt distribution therefore does not preclude either broad scenario of isolated or dynamical formation—or a mixture of both. However, the fact that the inferred minimum required tilt is significantly negative suggests the contribution of a dynamical formation channel to the BBH population. We further discuss the astrophysical interpretation of negative spin tilts and potential subpopulations in Section 6.3.2 in the context of effective spins.

6.3.2. *Effective Spins*

We next turn to the inferred distributions of the effective spins $\chi_{\rm eff}$ and $\chi_{\rm p}$. Figure 9 shows the marginal distributions of $\chi_{\rm eff}$ (left) and $\chi_{\rm p}$ (right) under the Skew-normal Effective Spin model (red solid; Equation B37), compared to the Gaussian Effective Spins model result from GWTC-3.0 (purple dashed; Equation B36). The Skew-normal Effective Spin model (Banagiri et al. 2025) differs from the previously-used Gaussian Effective Spins in two ways. First, it allows for asymmetry in the $\chi_{\rm eff}$ marginal distribution. Second, the marginal $\chi_{\rm p}$ distribution is a truncated Gaussian which is not correlated with $\chi_{\rm eff}$. On GWTC-4.0 data, the Gaussian Effective Spins model finds that $\chi_{\rm eff}$ and $\chi_{\rm p}$ are preferentially uncorrelated, although this conclusion depends on analysis settings; see Figure 19 in Appendix D.3 and the discussion therein.

We find that the $\chi_{\rm eff}$ distribution is skewed and asymmetric about zero with more support for positive values. Asymmetry can be directly probed with the ϵ skewness parameter of the Skew-normal Effective Spin model, defined in Equation (B37). We find that the skewness ϵ is less than 0 (symmetry) with 99.3% credibility, as can be seen in the inset of the left panel of Figure 9. This corresponds to a wider $\chi_{\rm eff}$ distribution to the right of the peak, and a narrower distribution to the left, and is known as *positive skew*. As discussed in Section 6.5.2, asymmetry is also found when allowing for a linear or spline correlation between $\chi_{\rm eff}$ and mass ratio. Results for the GWTC-4.0 $\chi_{\rm eff}$ distribution measured with more models, including the Gaussian Effective Spins model and with the weakly modeled approach are presented in Figures 19 and 21 in Appendix D.

We now discuss BBHs with spin tilts lying below the orbital plane. These interesting probes of dynamical BBH formation have $\cos\theta < 0$ and thus negative $\chi_{\rm eff}$. In Table 3, we present 90% bounds for three quantities which probe negative $\chi_{\rm eff}$ in the population, using both the Skew-normal Effective Spin and Gaussian Effective Spins models for a straightforward comparison to GWTC-3.0. First, we report the first percentile of $\chi_{\rm eff}$ distribution, which serves as a proxy for the distribution's minimum without necessitating the inclusion of sharp features in the distribution itself (Callister et al. 2022), constraining it to fall between -0.27 and -0.08. The Gaussian Effective Spins model finds more extremal minimum spins than the Skew-normal Effective Spin model, likely due to its required symmetry about its peak; fitting the positive side of the distribution forces a longer tail into the negative region. Second, we report the fraction of the population with negative $\chi_{\rm eff}$ to be 0.24–0.42, with the two models finding nearly identical results. Assuming isotropy, an upper bound on the fraction of BBHs formed dynamically in gas-free environments can be placed by doubling the fraction of negative $\chi_{\rm eff}$ (Equation 7 of Abbott et al. 2021a); we thus find that at most 84% of BBHs form dynamically. Third, we give the 90% upper limit on the fraction of BBHs coming from the hierarchical merger (HM) scenario. The HM fraction is bounded by the consideration that $\sim 16\%$ of BBHs coming from the HM formation channel will have $\chi_{\rm eff} < -0.3$ (Baibhav et al. 2020; Fishbach et al. 2022); we limit this fraction to $\lesssim 3\%$. Figure 21 in Appendix D.4 shows posterior distributions on the quantities in Table 3 for the models presented here and those using the weakly modeled approach.

It is possible that the support for $\chi_{\rm eff} < 0$ may be a byproduct of the models fitting for a peak at $\chi_{\rm eff} \sim 0$ without having the flexibility for a sharp decline beyond $\chi_{\rm eff} < 0$. This type of sharp feature could occur if component spin magnitudes cluster around $\chi \sim 0$ but their tilts seldom reach below $\cos\theta = 0$. In Section 6.3.1, we address this by directly fitting for the minimum required $\cos\theta$ which is found to be confidently negative, and corresponds to a fraction $0.41^{+0.05}_{-0.05}$ of systems with spins laying below the orbital plane. These probes of negative spin indicate that **between ∼20–40% the BBH population has spins which are more than 90 degrees misaligned with the orbital angular momentum**.

It is unlikely that the entire observed BBH population originates from a single formation channel (Zevin et al. 2021; Mandel & Broekgaarden 2022; Cheng et al.

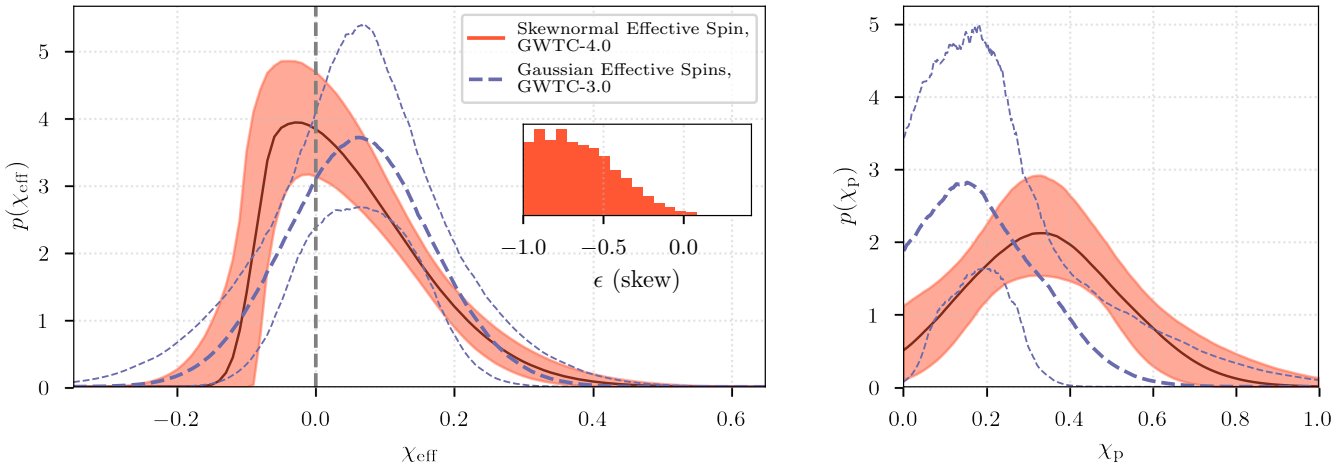

Figure 9. Marginal $\chi_{\rm eff}$ (left panel) and $\chi_{\rm p}$ (right panel) distributions. The solid lines show the median of each inferred distribution, and the shaded regions show the 90% credible intervals. The GWTC-4.0 results under the Skew-normal Effective Spin model are shown in red; the GWTC-3.0 results under the Gaussian Effective Spins model are shown in purple dashed. The histogram inset in the left panel shows the skew parameter ϵ for the Skew-normal Effective Spin model. There is significant preference for a skewed $\chi_{\rm eff}$ distribution, indicated by $\epsilon < 0$ with 99.3% credibility.

Table 3. Summary of probes of spin misalignment, as measured by $\chi_{\rm eff}$.

Model	$\chi_{\rm eff,1\%}$	Fraction $\chi_{\rm eff} < 0$	HM Fraction
Gaussian Effective Spins (GWTC-3.0)	$-0.18^{+0.09}_{-0.12}$	$0.28^{+0.12}_{-0.13}$	$< 3.1 \times 10^{-2}$
Gaussian Effective Spins (GWTC-4.0)	$-0.2^{+0.06}_{-0.07}$	$0.34^{+0.09}_{-0.1}$	$< 1.9 \times 10^{-2}$
Skew-normal Effective Spin (GWTC-4.0)	$-0.11^{+0.04}_{-0.08}$	$0.34^{+0.08}_{-0.09}$	$< 1.3 \times 10^{-4}$

Note— Summary of probes of spin misalignment, as measured by $\chi_{\rm eff}$. We give 90% credible intervals for the $\chi_{\rm eff}$ value of the first percentile of the distribution ($\chi_{\rm eff,1\%}$), serving as a proxy for the minimum $\chi_{\rm eff}$, and the fraction of $\chi_{\rm eff} < 0$. The HM fraction provides an upper limit to the fraction of BBHs of hierarchical merger origin, equal to 0.16 times the fraction of systems with $\chi_{\rm eff} < -0.3$ (Fishbach et al. 2022; Baibhav et al. 2020); we provide its 90% upper limit. Posteriors for these quantities for these and other models are plotted in Figure 21.

2023; Afroz & Mukherjee 2025; Colloms et al. 2025). Features in our spin distributions indeed suggest the presence of multiple sub-populations. A purely random spin channel would produce a symmetric distribution around $\chi_{\rm eff} = 0$ (Rodriguez et al. 2016a,b; Farr et al. 2017). However, for all models, the observed $\chi_{\rm eff}$ distribution is asymmetric about zero, c.f., Figure 21. This asymmetry suggests contribution from a preferentially aligned subpopulation (Gerosa et al. 2018; Arca Sedda et al. 2023; Banagiri et al. 2025). The skew observed in the Skew-normal Effective Spin, as well as the $(q, \chi_{\rm eff})$ Linear and $(q, \chi_{\rm eff})$ Spline models (see Section 6.5.2 and Figure 21) further supports the presence of an aligned component, either as a sub-dominant subpopulation or a dominant sub-population with small spin magnitudes. We find a preference for small spin magnitude (Figure 7), perhaps favoring the latter.

Finally, we discuss the effective precessing spin $\chi_{\rm p}$, which can provide additional insight about spin precession in the population. The right panel of Figure 9 shows the marginal inferred $\chi_{\rm p}$ distribution (red). We find that precession exists on a population level, as our models do not support $(\mu_p, \sigma_p) = (0, 0)$ at high credibility. The GWTC-4.0 results support larger $\chi_{\rm p}$ than GWTC-3.0, shown for comparison in Figure 9 (purple dashed). However, different likelihood convergence criteria (Appendix A.1) were used between GWTC-4.0 and GWTC-3.0, meaning these results are not directly comparable. Under the less-stringent GWTC-3.0 convergence criterion, the $\chi_{\rm p}$ distribution inferred from GWTC-4.0 data is more consistent with that found in

GWTC-3.0, albeit still with less support for low $\mu_{\rm p}$ and low $\sigma_{\rm p}$. The inferred $\chi_{\rm p}$ distribution is more dependent on analysis settings than other parameters, largely because the individual-event prior has no support at $\chi_{\rm p} = 0$ making it technically difficult to reweight individual-event posteriors to the population. We discuss the sensitivity of the $\chi_{\rm p}$ distribution to analysis settings in detail in Appendix D.3; see Figures 18 and 19.

6.4. *Merger Rate and Redshift Evolution*

Improvements in the sensitivity of current GW observatories (Ganapathy et al. 2023; Capote et al. 2025; Abac et al. 2025a) not only provide more BBH detections, but also observations of increasingly faint sources at higher redshifts. These observations allow us to improve our population-level constraints on the evolution of the merger rate across redshift.

Following Abbott et al. (2023a), we repeat the Power Law Redshift strongly modeled approach. We assume that the merger rate evolves as $\mathcal{R}(z) \propto (1+z)^{\kappa}$, and we infer the proportionality constant and the power-law index κ. **We find that the BBH merger rate at $z = 0.2$ is $29^{+8.5}_{-6.5}\,\mathrm{Gpc^{-3}\,yr^{-1}}$, and the power-law exponent is constrained to $\kappa = 3.2^{+0.94}_{-1.00}$,** which represent consistent and improved constraints over our GWTC-3.0 analysis. We infer that 99% of detectable BBHs fall below $z = 1.5^{+0.2}_{-0.2}$ (cf. the maximum observable redshift inferred as $z = 1.1^{+0.2}_{-0.2}$ in Fishbach & van Son 2023, from GWTC-3.0). We also use the weakly modeled B-Spline approach to explore if the data support behavior beyond a power-law evolution, and compare our results in Figure 10. While the results are consistent—indicating that the Power Law Redshift model is sufficient—the B-Spline approach infers a larger merger rate of $38^{+19}_{-10}\,\mathrm{Gpc^{-3}\,yr^{-1}}$ at $z = 0.2$. Our results are nominally consistent with the cosmic star formation rate density, with $\kappa_{\rm SFR} = 2.7$ (Madau & Dickinson 2014).

Our constraints on the merger rate evolution informs our understanding of the BBH progenitor formation rate and the delay time distribution between formation and merger (Vitale et al. 2019; Rodriguez & Loeb 2018; Fragione & Kocsis 2018; Fishbach et al. 2018; Baibhav et al. 2019; Romero-Shaw et al. 2021; Broekgaarden et al. 2022b; Mapelli et al. 2022; Fishbach & Kalogera 2021; Chruślińska 2024; Fishbach & Fragione 2023; Boesky et al. 2024). For BBHs formed from isolated binary evolution, the merger rate evolution is typically well approximated as a power law at small redshift, though the value of the power-law index κ is sensitive to the assumed population synthesis parameters (Neijssel et al. 2019; Broekgaarden et al. 2022b; Gallegos-Garcia et al. 2021; de Sá et al. 2024). Nevertheless, models typically prefer values around $\kappa \sim 1$ (Dominik et al. 2013; Baibhav et al. 2019). BBHs originating from dense star clusters predict local merger rates of $\mathcal{R} \sim 10\,\mathrm{Gpc^{-3}\,yr^{-1}}$ (Arca Sedda et al. 2024) and steeper values of $\kappa \sim 2$ (Antonini & Gieles 2020), albeit with large theoretical uncertainties that can account for a similar evolution to the isolated binary evolution predictions. **Our observations suggest a steeper evolution, with $\kappa \sim 2$–4.** While this does not rule out either class of BBH formation, it may be suggestive of (i) progenitor formation rates that peak at an earlier redshift as compared to the cosmic star formation rate density, e.g., due to a preference for low-metallicity progenitors, and/or (ii) shorter delay times, perhaps with a tail toward long delay times (Fishbach

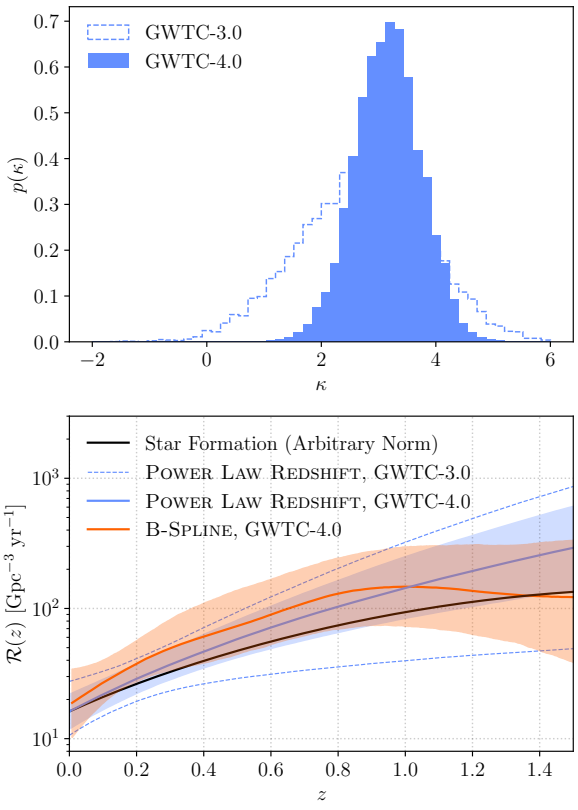

Figure 10. Comparison of redshift models between GWTC-3.0 and GWTC-4.0. *Top:* Posterior on the κ parameter for the Power Law Redshift model. GWTC-4.0 shows increased support for positive values. *Bottom:* Median and 90% credible regions for the comoving source frame rate in the Power Law Redshift model. We also show a comparison of the Power Law Redshift model and the weakly modeled B-Spline model and a scaled cosmic star formation rate density (Madau & Dickinson 2014), which are consistent within uncertainties ($\kappa_{\rm SFR} = 2.7$).

& Kalogera 2021; Karathanasis et al. 2023; Fishbach & van Son 2023; Turbang et al. 2024; Vijaykumar et al. 2024; Schiebelbein-Zwack & Fishbach 2024).

Models of isolated binary evolution and dense star clusters often predict that the merger rate evolves differently across the mass spectrum (van Son et al. 2022b; Mapelli et al. 2022; Ye & Fishbach 2024); however, analyses of the previous catalog have not found evidence for or against differential rate evolution (Fishbach et al. 2021; Sadiq et al. 2022; van Son et al. 2022b; Ray et al. 2023a; Heinzel et al. 2025b; Sadiq et al. 2025a). We discuss potential mass-redshift correlations in Section 6.5.4.

6.5. *Population-level Correlations between Parameters*

While much can be learned by studying the distributions of individual BBH parameters, we can glean additional information on the population by considering how parameters are correlated with one another across systems. In this subsection, we provide an overview of how parameters in the BBH population appear to be broadly structured in various two-dimensional slices of parameter space, briefly discussing the astrophysical implications of our findings.

6.5.1. *Mass Ratio and Spin Correlations*

We begin by following up on the purported anti-correlation between BBH mass ratio q and effective inspiral spin $\chi_{\rm eff}$. **Although we find less support for the specific case of anti-correlation between q and $\chi_{\rm eff}$ relative to GWTC-3.0, we find compelling evidence for some correlated structure in $(q, \chi_{\rm eff})$.** Namely, larger positive values of $\chi_{\rm eff}$ appear to be favored as q decreases, but it is unclear if this is accompanied by a preference for larger negative values of $\chi_{\rm eff}$ as well. The $(q, \chi_{\rm eff})$ Linear model imposes a linear functional dependence between BBH mass ratio, and the mean and (natural log) width of the $\chi_{\rm eff}$ distribution (see Appendix B.7). Fitting this model to GWTC-3.0 data suggests that a linear correlation coefficient of $\delta\mu_{\rm eff|q} < 0$ (i.e., the case of a negative q-dependence on the mean of the $\chi_{\rm eff}$ distribution) with 98% credibility (Abbott et al. 2023a). Updating this analysis to include data obtained over O4a softens the evidence for an anti-correlation between q and the mean of the $\chi_{\rm eff}$ distribution, with a value of $\delta\mu_{\rm eff|q} < 0$ now inferred at 82% credibility. However, we now see notable evidence for a linear increase in the log width of the $\chi_{\rm eff}$ distribution as mass ratios become more unequal – with $\delta \ln \sigma_{\rm eff|q} < 0$ at 95% credibility.

We plot the mass-ratio dependent mean and width of the $\chi_{\rm eff}$ distribution inferred using the Linear model in Figure 11. In the bottom panel of Figure 11, we include the two-dimensional posterior distribution of $\delta\mu_{\rm eff|q}$ and $\delta \ln \sigma_{\rm eff|q}$. Inspecting this plot, we see an anti-correlation in the posterior of the two hyperparameters, where larger negative values of $\delta\mu_{\rm eff|q}$ imply values of $\delta \ln \sigma_{\rm eff|q}$ closer to zero, and larger negative values of $\delta \ln \sigma_{\rm eff|q}$ imply values of $\delta\mu_{\rm eff|q}$ closer to zero. The case in which both parameters are zero (no correlation between q and $\chi_{\rm eff}$ of any kind) appears to be ruled out at $> 99\%$ credibility. In practice, this implies that the Linear model finds support for larger positive values of $\chi_{\rm eff}$ at more unequal mass ratios, but cannot yet conclude whether these are accompanied by larger negative values of $\chi_{\rm eff}$ as well.

Next, we probe for more intricate correlations between q and $\chi_{\rm eff}$ using the Spline model. Similar to the Linear model, the Spline model allows for the mean and width of the $\chi_{\rm eff}$ distribution to evolve with mass ratio. However, these mass ratio dependences, rather than being linear, are modeled flexibly with cubic splines (see Appendix B.7 and Heinzel et al. 2024). We plot the q-dependent means and widths inferred from the Spline model alongside those from the Linear model in Figure 11. Despite fluctuations emerging in the Spline model, the two models are, within 90% credible bounds, consistent. The preference for broadening in the $\chi_{\rm eff}$ distribution as q decreases found with the Linear model, does not clearly appear in the Spline model. In the bottom panel of Figure 11, we also plot the inferred gradient of the $\chi_{\rm eff}$ distribution's mean and (natural log) width relative to mass ratio at $q = 0.6$ (roughly the value at which covariance appears most pronounced). Here, we see a similar structure to that of the Linear model's $(\delta\mu_{\rm eff|q}, \delta \ln \sigma_{\rm eff|q})$ posterior, albeit with much more uncertainty. As a final observation from the Spline model in Figure 11, we see that the inferred $\chi_{\rm eff}$ distribution at low mass ratios ($q \lesssim 0.2$) grows in uncertainty, roughly recovering the prior. This implies that any correlation inferred by the Linear model is likely driven by observations with mass ratios $q \gtrsim 0.2$. As such, the inferences from this model should be interpreted with caution as $q \to 0$, where the trend is effectively being extrapolated from the more populated $q \gtrsim 0.2$ region due to a lack of model flexibility.

We now move to the Copula model, which allows for a correlation between q and $\chi_{\rm eff}$ with a variable strength $\kappa_{q,\rm eff}$ (see Appendix B.6). This framework has the added advantage that the level of correlation is decoupled from the shape of the marginal distribution (e.g., Adamcewicz & Thrane 2022), with less flexibility for covariant structure relative to other correlated population models. Fitting the Copula model to GWTC-4.0 suggests that q and $\chi_{\rm eff}$ are anti-correlated ($\kappa_{q,\rm eff} < 0$) with 92% cred-

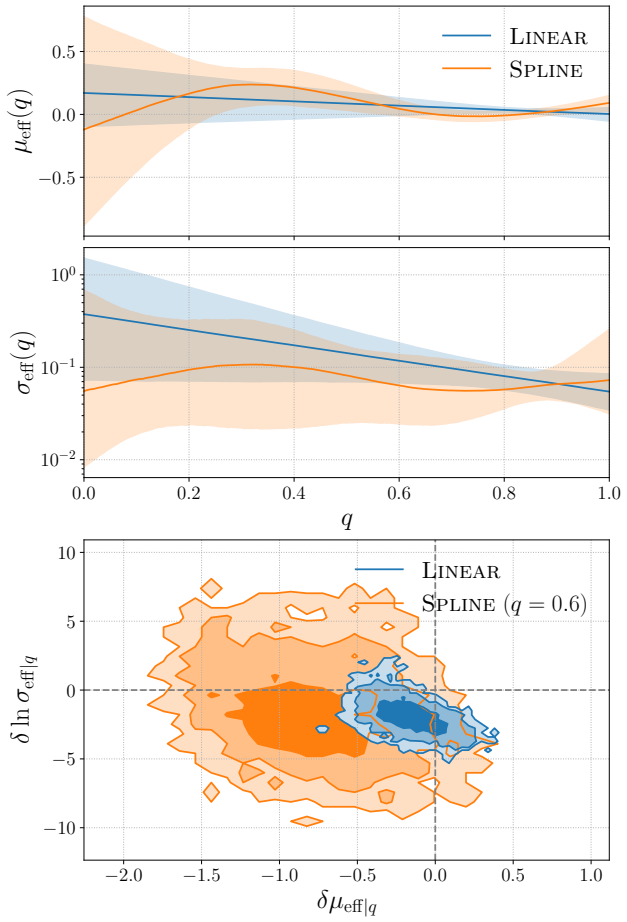

Figure 11. The inferred peak (top) and width (middle) of the $\chi_{\rm eff}$ distribution as a function of mass ratio for the $(q, \chi_{\rm eff})$ Linear model (blue) and Spline model (orange). The shaded regions in these panels give the 90% credible intervals. The bottom panel gives the posterior distribution for the gradient of the $\chi_{\rm eff}$ distribution's peak ($\delta\mu_{{\rm eff}|q}$) and natural log width ($\delta \ln \sigma_{{\rm eff}|q}$) dependent on mass ratio. From dark to light, the shaded regions represent the 50%, 90% and 99% credible intervals. Again, blue gives the result of the Linear model, while orange shows the result of the Spline model sliced through $q = 0.6$ (the approximate point at which the gradients are largest). It appears that mass ratio and $\chi_{\rm eff}$ exhibit some kind of correlation, but the exact nature is unclear.

ibility. Specifically, we infer $\kappa_{q,{\rm eff}} = -2.1^{+2.4}_{-2.9}$. Adamcewicz et al. (2023) analyzed GWTC-3.0 data with a copula model to find that $(q, \chi_{\rm eff})$ are anti-correlated with $> 99\%$ credibility, suggesting greater evidence for an anti-correlation than is measured here. However, these results are not directly comparable to those presented in this work, due to different modeling assumptions and different convergence criteria in the population likelihood. Qualitatively, we see that the Copula model and Linear model exhibit a subtle anti-correlation in

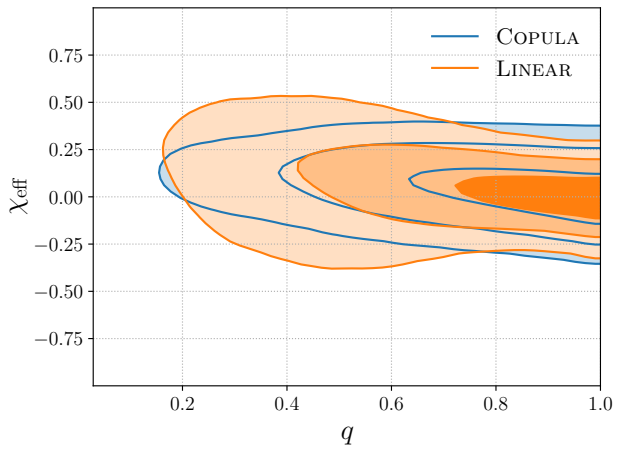

Figure 12. Mass ratio and $\chi_{\rm eff}$ PPDs for the Copula model (blue) and Linear model (orange). The contours, from dark to light, mark 50%, 90%, and 99% of the volume. We see a subtle preference for an anti-correlation, although this feature is far less prevalent than it appeared in GWTC-3.0. We also see evidence for a broadening in the distribution as mass ratios become more unequal in the Linear model. The Copula model is not flexible in a way that it can capture a broadening.

$(q, \chi_{\rm eff})$, as seen in the respective two-dimensional PPDs in Figure 12.

There are a number of potential astrophysical implications if the anti-correlation in $(q, \chi_{\rm eff})$ is real. If isolated binaries make up a substantial fraction of the population and undergo tidal spin up, stable mass transfer and the resulting mass ratio reversal of systems (Broekgaarden et al. 2022a; Zevin & Bavera 2022; Olejak et al. 2024) may produce an such an anti-correlation in the population. This feature could also be explained by binaries undergoing a common-envelope phase provided common envelope efficiencies are sufficiently high (Bavera et al. 2021).

Covariance in $(q, \chi_{\rm eff})$ could also be a result of hierarchical mergers contributing to a considerable fraction of the BBH merger rate (Antonini et al. 2025b,a). Hierarchical mergers should exhibit mass ratios that are more unequal and spin magnitudes that are larger than isolated, or first-generation dynamical mergers. If hierarchical mergers occur in typical dynamical environments, the spins of the BHs will be isotropically distributed, thus producing a broadening in the $\chi_{\rm eff}$ distribution as mass ratios become unequal. Meanwhile, BHs in AGN may have spins preferentially aligned with one another, meaning hierarchical mergers in these environments can produce a correlation between mass ratios and spins that is asymmetrical about $\chi_{\rm eff} = 0$ (McKernan et al. 2022; Santini et al. 2023). Stronger evidence for unequal mass binaries preferring positive values of $\chi_{\rm eff}$ could then in-

dicate that binaries merging in AGN are predominantly coaligned with the rotation of the AGN disk (Santini et al. 2023).

6.5.2. *Mass and Spin Correlations*

We find model-dependent evidence for correlations between m_1 and $\chi_{\rm eff}$. Motivated by the feature in the mass distribution around the $\sim$30$-$40 $M_\odot$ range, we use a BGP analysis allowing for correlations in m_1 and $\chi_{\rm eff}$ (see Appendix C.2) to see if this feature in the distribution of masses is accompanied by a deviation in the BBH spin distribution. Using GWTC-4.0, these studies find weak evidence that BBH systems with at least one mass in the $\sim$30$-$40 $M_\odot$ peak tend to have spins that are symmetrically distributed about $\chi_{\rm eff} = 0$, while binaries outside of this mass range have spins that are skewed toward positive (aligned) values of $\chi_{\rm eff}$. Quantitatively, this BGP analysis infers that for every merger in the $\sim$30$-$40 $M_\odot$ peak with a negative value of $\chi_{\rm eff}$, there are $1.9^{+4.8}_{-1.2}$ with a positive value of $\chi_{\rm eff}$. Meanwhile, for every event outside of this mass range with a negative value of $\chi_{\rm eff}$, there are $6.6^{+14.1}_{-4.2}$ with a positive value of $\chi_{\rm eff}$. We illustrate this result in the top panel of Figure 13. Furthermore, in the BGP analysis the inferred distribution of masses for BBH systems with effective inspiral spins ($|\chi_{\rm eff}| \lesssim 0.1$) exhibits a preference for a larger proportion of mergers with $m_1 \sim 30$$-$$40\,M_\odot$, compared to systems with $\chi_{\rm eff} \gtrsim 0.1$. These distributions, however, remain consistent within 90% credible intervals. This is illustrated in the bottom two panels of Figure 13. These features were also recovered with same analyses applied to GWTC-3.0 data, albeit with less certainty (Ray et al. 2024). However, more recent GWTC-3.0 analyses find conflicting evidence, suggesting that the $\chi_{\rm eff}$ distribution of $\approx 30\,M_\odot$ BHs is positively skewed (Sadiq et al. 2025b; Roy et al. 2025).

The Isolated Peak model (see also Godfrey et al. 2023) fits the data to a model containing multiple mass subpopulations with independent spin distributions. One subpopulation consists of a single peak in the mass distribution (which is inferred to center on $\sim$10 $M_\odot$), while other subpopulations are flexibly fit with the B-Spline method. Applying these analyses to GWTC-4.0 suggests that BHs within the $\sim$10 $M_\odot$ peak have a spin magnitude distribution consistent (within 90% credibile intervals) with the rest of the BBH population. Meanwhile, although there is some overlap in the 90% credibile regions of the cosine tilt distributions for both subpopulations, the distribution of tilts for BHs in the $\sim$10 $M_\odot$ peak favors alignment and is inconsistent with isotropy. The distribution of cosine tilts for BHs outside this mass peak appears more symmetrical

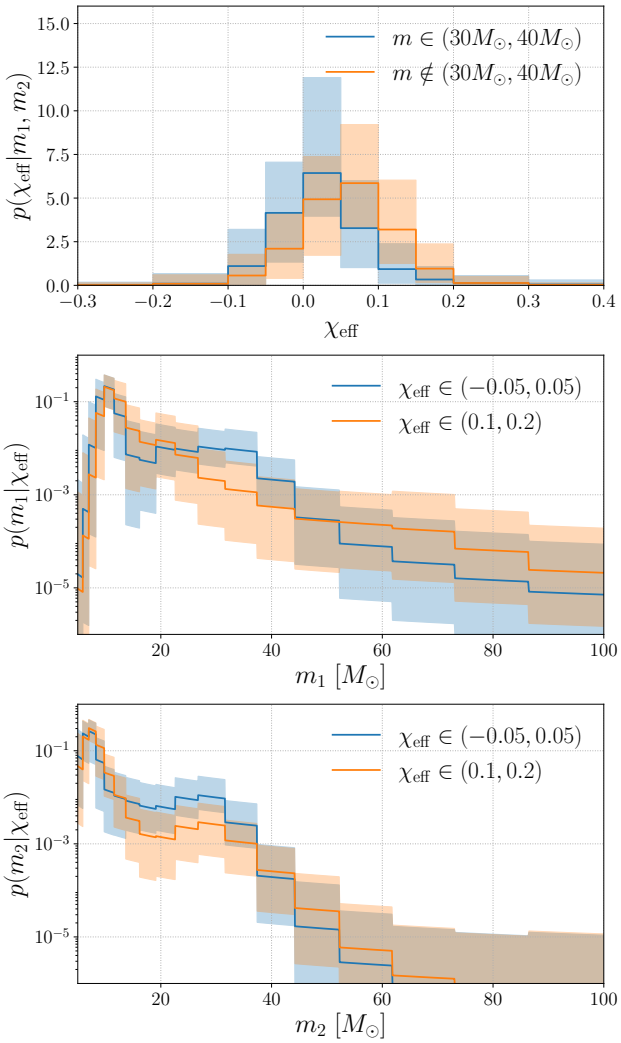

Figure 13. *Top:* inferred distributions of effective inspiral spin $\chi_{\rm eff}$ in the correlated mass–spin BGP analysis. The solid lines give the median, while the shaded regions indicate the 90% credible intervals. In blue, we have the $\chi_{\rm eff}$ distribution of BBH systems with at least one mass inside the 30$-$40 $M_\odot$ peak range. In orange, we have the $\chi_{\rm eff}$ distribution of BBH systems masses outside the 30$-$40 $M_\odot$ peak range. While the two distributions are consistent within 90% credible intervals, systems outside the peak range appear to favour a distribution of $\chi_{\rm eff}$ skewed toward more positive values. Meanwhile, systems inside the peak range appear to favour a more symmetrical distribution about $\chi_{\rm eff} = 0$. *Middle:* inferred distributions of primary mass from the correlated mass–spin BGP analysis. *Bottom:* inferred distributions of secondary mass from the correlated mass–spin BGP analysis. In both mass plots, blue shows the distribution for systems with small spins, $\chi_{\rm eff}$ in the range $(-0.05, 0.05)$, while orange shows the distribution for systems with larger spins, $\chi_{\rm eff}$ in the range $(0.1, 0.2)$. Note the preference for a higher proportion of mergers with masses $\approx$ 30$-$40 $M_\odot$ for low spin systems (blue).

about $\cos\theta = 0$ and does not rule out isotropy. Roughly speaking, this corresponds to a distribution of $\chi_{\rm eff}$ that may be symmetrical about zero for the BBH population outside of the $\sim$10 $M_\odot$ peak. Inside this mass peak, however, the implied $\chi_{\rm eff}$ distribution skews toward positive values.

We also model a correlation between primary mass and $\chi_{\rm eff}$ using a copula model (see Appendix B.6). Fitting this Copula model with GWTC-4.0 data, we infer a correlation of $\kappa_{m_1,\rm eff} = 0.4^{+1.6}_{-1.4}$ between m_1 and $\chi_{\rm eff}$. Hence, the Copula analysis finds no evidence for a single, smooth, population-wide correlation between primary mass and $\chi_{\rm eff}$.

Broadly, our analyses of the joint mass and spin distribution recover similar features to recent works that probe for features in GWTC-3.0. First, analyses of GWTC-3.0 do not find evidence for a smoothly correlated distribution in primary mass and $\chi_{\rm eff}$ (Safarzadeh et al. 2020; Biscoveanu et al. 2022a; Fishbach et al. 2022; Heinzel et al. 2024, 2025b; Antonini et al. 2025b), as is the case in this work. Meanwhile, a number of analyses find evidence for separate subpopulations in mass and spin consistent with predictions of dynamical and field mergers. More specifically, these works tend to find a subpopulation with lower masses and smaller, mostly aligned spins, and a second subpopulation of heavier binaries with larger, isotropically distributed spins, where the subpopulations transition around $\sim$40 $M_\odot$ (Wang et al. 2022; Mould et al. 2022b; Godfrey et al. 2023; Li et al. 2024; Antonini et al. 2025b; Pierra et al. 2024; Guo et al. 2024; Li et al. 2025; Sadiq et al. 2025b). While the analyses presented in this work do not recover this exact feature (nor do they probe directly for it), the Isolated Peak model's preference for aligned spins at $\sim$10 $M_\odot$ and isotropically distributed spins at higher masses may be related.

6.5.3. *Redshift and Spin Correlations*

We find evidence that the $\chi_{\rm eff}$ distribution broadens as redshift increases up to $z \sim 1$. We again look for correlations between redshift and $\chi_{\rm eff}$ by modeling the mean and width of the $\chi_{\rm eff}$ distribution with a linear dependence on redshift (see Appendix B.7). Biscoveanu et al. (2022a) employed a Linear model for $(z, \chi_{\rm eff})$ to analyze GWTC-3.0 data, finding no evidence for redshift dependence in the mean of the $\chi_{\rm eff}$ distribution, but suggesting that the $\chi_{\rm eff}$ distribution broadens with increasing redshift ($\delta \ln \sigma_{\rm eff|z} > 0$ at 99% credibility). Updating this analysis to include O4a data, we find more evidence for a broadening in the $\chi_{\rm eff}$ distribution with redshift, with $\delta \ln \sigma_{\rm eff|z} > 0$ now inferred at $> 99\%$ credibility. The mean and width of the $\chi_{\rm eff}$ distribution as functions of redshift are shown in Figure 14.

The Spline model, which models the redshift dependence on the mean and width of the $\chi_{\rm eff}$ distribution flexibly with cubic splines (see Appendix B.8), is broadly consistent with the Linear model. The exception is that the Spline model tends to recover the prior beyond a redshift of $z \gtrsim 1$. This likely indicates that the Linear model is fitting a trend at low redshifts, then extending this trend to redshifts $z \gtrsim 1$ due to a lack of flexibility. To further illustrate this point, the gray dashed line in Figure 14 indicates the redshift under which 90% of the catalog's cumulative posterior support is contained. Roughly speaking, this means our inferences above this redshift are informed by only $\sim$10% of the data. Therefore, while we find evidence that the $\chi_{\rm eff}$ distribution broadens as binaries approach $z \sim 1$, we are unable to determine if this trend continues at higher redshifts.

We also model a correlation between redshift and $\chi_{\rm eff}$ using a copula model, with variable correlation $\kappa_{z,\rm eff}$ (see Appendix B.6). In contrast to the above analyses, the Copula model finds evidence for a positive correlation in $(z, \chi_{\rm eff})$, inferring a value of $\kappa_{z,\rm eff} = 4.6^{+3.0}_{-3.1}$ (or $\kappa_{z,\rm eff} > 0$ with 98% credibility). While the Copula model lacks the flexibility to fit a broadening directly, the long (albeit shallow) posterior tails reaching into both large-negative and large-positive values for $\kappa_{z,\rm eff}$ may relate to the broadening in the $\chi_{\rm eff}$ distribution recovered by the Linear and Spline models (see the subtle mode at negative values of $\kappa_{z,\rm eff}$ in Appendix D.6). If the z-dependent broadening and constant mean of the $\chi_{\rm eff}$ distribution from the Linear and Spline models is to be believed, it is unclear why the posterior on $\kappa_{z,\rm eff}$ skews toward positive values, rather than being symmetric about zero.

It is unclear what this broadening in the $\chi_{\rm eff}$ distribution at greater redshifts might imply astrophysically. Provided BH progenitors experience significant spin up due to tidal torques, smaller orbital separations and periods will result in more efficient spin up of BBH components (Zaldarriaga et al. 2018; Mapelli 2020; Bavera et al. 2021, 2022; Fuller & Lu 2022). Of course, smaller orbital separations correspond to shorter delay-times to merger. This means that these systems may contribute more to the merger rate in the earlier Universe (higher z), where systems with wider orbits have not had time to merge. Furthermore, given that higher metallicity systems are expected to evolve to longer orbital periods due to experiencing more mass loss during core helium burning (Qin et al. 2018; Fuller & Lu 2022), one might predict that larger spin magnitudes should be ob-

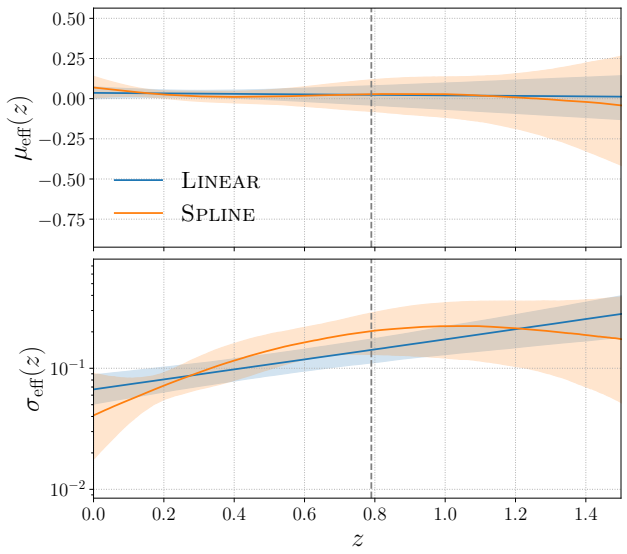

Figure 14. The inferred mean (top) and width (bottom) of the $\chi_{\rm eff}$ distribution as a function of redshift for the LINEAR model (blue) and the SPLINE model (orange). The shaded regions give the 90% credible intervals. The gray dashed line indicates the redshift under which 90% of the cumulative posterior probability lies across all events in GWTC-4.0. Therefore, inferences above this redshift are dominated by the prior or population model. Note the logarithmic scale on the width (bottom panel). In either model, we find evidence that the width of the effective inspiral spin distribution increases with redshift up to $z \sim 1$.

served in regions of lower metallicity and thus also at higher redshifts. However, the correlation in the LINEAR and SPLINE models is between redshift and the width of the $\chi_{\rm eff}$ distribution rather than the mean. This includes an increasing number of systems with both large-positive and large-negative values of $\chi_{\rm eff}$. An increase in systems with large-negative $\chi_{\rm eff}$ may be somewhat difficult to square with the above hypothesis, given that tidal spin up is only relevant in isolated systems, and that large supernovae kicks are required to misalign the BH spins so significantly (Kalogera 2000; Wysocki et al. 2018; Gerosa et al. 2018; Callister et al. 2021a; Steinle & Kesden 2021; Stevenson 2022; Tauris 2022; Baibhav & Kalogera 2024). On the other hand, if the COPULA results are to be taken at face-value, the preference for a positive correlation in $(z, \chi_{\rm eff})$ fits more neatly with the aforementioned hypothesis.

Another possibility, following discussion from Section 6.3, is that multiple separate subpopulations of BBH mergers are being observed, where the more dominant (in terms of redshift-dependent merger rate) subpopulation changes at some nearby redshift. Hierarchical mergers becoming more dominant at higher redshifts for example could potentially explain such a broadening in the effective spin distribution. This interpretation would also imply some level of correlation between mass and redshift, which we do not find evidence for below in Section 6.5.4.

As a final caveat, using GWTC-3.0 data, Biscoveanu et al. (2022a) fit the BBH population to a model in which the $\chi_{\rm eff}$ distribution is linearly dependent on both primary mass and redshift. In doing so, the authors find degeneracies between the $(m_1, \chi_{\rm eff})$ and $(z, \chi_{\rm eff})$ correlations. While we do not explore correlations beyond two dimensions in this work, we acknowledge that the assumption of independence between other pairs of parameters may have notable effects on our inferences.

6.5.4. *Redshift and Mass Correlations*

Finally, we consider potential correlations between BBH mass and redshift. **We do not find evidence for evolution in the mass distribution with redshift.** We emphasize that these inferences are constrained to the nearby Universe, with relatively little data beyond $z \sim 1$ (19 out of 153 BBH events having posteriors consistent with $z > 1$ to 90% credibility). We model correlations between primary mass and redshift using a copula with correlation $\kappa_{m_1,z}$ (see Appendix B.6). Using this COPULA model, we infer a correlation of $\kappa_{m_1,z} = 0.6^{+2.8}_{-2.6}$. We also search for correlations between mass and redshift using a BGP analysis that allows for covariance in m_1 and z (see Appendix C.2). Similarly, this analysis finds that the mass distribution does not show a distinguishable evolution with redshift (see Appendix D.6). This conclusion is mostly in line with studies using GWTC-3.0, which are also unable to infer any correlation between mass and redshift with confidence (Fishbach et al. 2021; Sadiq et al. 2022; van Son et al. 2022b; Karathanasis et al. 2023; Ray et al. 2023a; Heinzel et al. 2024, 2025b; Lalleman et al. 2025; Sadiq et al. 2025a). This is with the exception of Rinaldi et al. (2024) finding a positive correlation between BBH primary mass and redshift, although a novel method is used to account for selection effects.

7. CONCLUSION

In this paper, we present population-level analyses of events included in the fourth Gravitational-Wave Transient Catalog GWTC-4.0. This dataset more than doubles the number of events analyzed compared to the previous catalog. Our main findings are:

1. Features identified in the third catalog GWTC-3.0 persist, including clear overabundances in the mass distribution at 1–2 $M_\odot$ and around 10 $M_\odot$, and a feature near 35 $M_\odot$. There is no conclusive evi-

dence to either support or refute a suppression of the merger rate between these features.

2. We estimate the merger rates at redshift $z = 0$ to be 7.6–250 $\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ for binary neutron stars, 9.1–84 $\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ for neutron star–black hole binaries, 14–26 $\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$ for binary black holes.

3. The binary black hole primary mass distribution is well described by a broken power law, shallow at low masses and steep at high masses, modulated by overdensities near $10\,M_\odot$ and $35\,M_\odot$. A weakly modeled approach finds evidence of overdensity around $20\,M_\odot$.

4. Black holes in the $35\,M_\odot$ feature tend to pair with companions of similar mass more frequently than lower-mass black holes do.

5. The distribution of effective inspiral spins is asymmetric about $\chi_{\mathrm{eff}} = 0$ and is skewed toward positive χ_{eff} values. Spin magnitudes span a broad range from 0 to 1, although ∼90% of BHs have $\chi < 0.57$.

6. The redshift evolution of the binary black hole merger rate $\mathcal{R}(z)$ remains consistent with the cosmic star formation rate density. A merger rate uniform in comoving volume and source-frame time is ruled out.

7. We find that black holes outside the 30–40 $M_\odot$ range prefer an asymmetric effective inspiral spin distribution skewed toward higher values, while those within this range show no such preference. Compared to GWTC-3.0, we observe stronger evidence that the width of the effective spin distribution increases with redshift and weaker evidence for an anti-correlation between mass ratio and effective spin. No redshift evolution is observed in the mass distribution.

8. The neutron star mass distribution remain consistent with previous results, favoring a broad distribution of neutron star masses between $1\,M_\odot$ and $3\,M_\odot$.

The analysis of this expanded dataset refines the statistical significance of previously reported trends and reveals new features in the population of compact binary mergers. Several of these findings, such as the persistent $\sim 10\,M_\odot$ feature, pose challenges to current models of supernova physics, binary mass transfer, and dynamical formation in dense stellar environments. Future data from the remainder of the fourth observing run O4 will further enhance our understanding and may uncover additional structure in the population.

ACKNOWLEDGMENTS

This material is based upon work supported by NSF's LIGO Laboratory, which is a major facility fully funded by the National Science Foundation. The authors also gratefully acknowledge the support of the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO 600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. The authors gratefully acknowledge the Italian Istituto Nazionale di Fisica Nucleare (INFN), the French Centre National de la Recherche Scientifique (CNRS) and the Netherlands Organization for Scientific Research (NWO) for the construction and operation of the Virgo detector and the creation and support of the EGO consortium. The authors also gratefully acknowledge research support from these agencies as well as by the Council of Scientific and Industrial Research of India, the Department of Science and Technology, India, the Science & Engineering Research Board (SERB), India, the Ministry of Human Resource Development, India, the Spanish Agencia Estatal de Investigación (AEI), the Spanish Ministerio de Ciencia, Innovación y Universidades, the European Union NextGenerationEU/PRTR (PRTR-C17.I1), the ICSC - CentroNazionale di Ricerca in High Performance Computing, Big Data and Quantum Computing, funded by the European Union NextGenerationEU, the Comunitat Autonòma de les Illes Balears through the Conselleria d'Educació i Universitats, the Conselleria d'Innovació, Universitats, Ciència i Societat Digital de la Generalitat Valenciana and the CERCA Programme Generalitat de Catalunya, Spain, the Polish National Agency for Academic Exchange, the National Science Centre of Poland and the European Union - European Regional Development Fund; the Foundation for Polish Science (FNP), the Polish Ministry of Science and Higher Education, the Swiss National Science Foundation (SNSF), the Russian Science Foundation, the European Commission, the European Social Funds (ESF), the European Regional Development Funds (ERDF), the Royal Society, the Scottish Funding Council, the Scottish Universities Physics Alliance, the Hungarian Scientific Research Fund (OTKA), the French Lyon Institute of Origins (LIO), the Belgian Fonds de la Recherche Scientifique (FRS-FNRS), Actions de Recherche Concertées (ARC) and Fonds Wetenschappelijk Onderzoek - Vlaan-

deren (FWO), Belgium, the Paris Île-de-France Region, the National Research, Development and Innovation Office of Hungary (NKFIH), the National Research Foundation of Korea, the Natural Sciences and Engineering Research Council of Canada (NSERC), the Canadian Foundation for Innovation (CFI), the Brazilian Ministry of Science, Technology, and Innovations, the International Center for Theoretical Physics South American Institute for Fundamental Research (ICTP-SAIFR), the Research Grants Council of Hong Kong, the National Natural Science Foundation of China (NSFC), the Israel Science Foundation (ISF), the US-Israel Binational Science Fund (BSF), the Leverhulme Trust, the Research Corporation, the National Science and Technology Council (NSTC), Taiwan, the United States Department of Energy, and the Kavli Foundation. The authors gratefully acknowledge the support of the NSF, STFC, INFN and CNRS for provision of computational resources.

This work was supported by MEXT, the JSPS Leading-edge Research Infrastructure Program, JSPS Grant-in-Aid for Specially Promoted Research 26000005, JSPS Grant-in-Aid for Scientific Research on Innovative Areas 2402: 24103006, 24103005, and 2905: JP17H06358, JP17H06361 and JP17H06364, JSPS Core-to-Core Program A. Advanced Research Networks, JSPS Grants-in-Aid for Scientific Research (S) 17H06133 and 20H05639, JSPS Grant-in-Aid for Transformative Research Areas (A) 20A203: JP20H05854, the joint research program of the Institute for Cosmic Ray Research, University of Tokyo, the National Research Foundation (NRF), the Computing Infrastructure Project of the Global Science experimental Data hub Center (GSDC) at KISTI, the Korea Astronomy and Space Science Institute (KASI), the Ministry of Science and ICT (MSIT) in Korea, Academia Sinica (AS), the AS Grid Center (ASGC) and the National Science and Technology Council (NSTC) in Taiwan under grants including the Science Vanguard Research Program, the Advanced Technology Center (ATC) of NAOJ, and the Mechanical Engineering Center of KEK.

Additional acknowledgements for support of individual authors may be found in the following document: `https://dcc.ligo.org/LIGO-M2300033/public`.

APPENDIX

A. HIERARCHICAL INFERENCE DETAILS

A.1. *Likelihood Estimation*

The analytic integrals in Equation (2) are not tractable, and so we estimate the integrals using Monte Carlo estimation (Tiwari 2018; Farr 2019; Essick & Farr 2022; Talbot & Golomb 2023). For example, the estimator for the likelihood $\hat{\mathcal{L}}(d_i|\mathbf{\Lambda})$ is

$$\hat{\mathcal{L}}(d_i|\mathbf{\Lambda}) \propto \frac{1}{N_{\rm PE}} \sum_{j=1}^{N_{\rm PE}} \frac{\pi(\boldsymbol{\theta}_{ij}|\mathbf{\Lambda})}{p(\boldsymbol{\theta}_{ij})}, \tag{A1}$$

where $\{\boldsymbol{\theta}_{ij}\}_{j=1}^{N_{\rm PE}}$ are a collection of $N_{\rm PE}$ samples from the posterior on the GW parameters of the ith event d_i. We divide out by the parameter estimation prior $p(\boldsymbol{\theta})$, and so the Monte Carlo estimator in Equation (A1) converges to the ith integral inside the product of Equation (2) in the limit of $N_{\rm PE} \to \infty$.

Similarly, the estimator for the selection efficiency $\hat{\xi}$ is (Tiwari 2018; Farr 2019; Essick & Farr 2022)

$$\hat{\xi}(\mathbf{\Lambda}) \propto \frac{1}{N_{\rm draw}} \sum_{j=1}^{N_{\rm found}} \frac{\pi(\boldsymbol{\theta}_j|\mathbf{\Lambda})}{\pi(\boldsymbol{\theta}_j|\mathbf{\Lambda}_{\rm draw})}, \tag{A2}$$

where $N_{\rm draw}$ events with parameters $\boldsymbol{\theta}_i$ are injected into representative noise from the detectors. The search pipelines then search these synthetically generated data and recover some subset $N_{\rm found}$ of the events with the detection statistic exceeding some threshold. In the limit of $N_{\rm draw} \to \infty$, this approaches the true integral in Equation (3).

Because we only have a finite number of samples from each event and finite $N_{\rm draw}$, we must be careful to account for the intrinsic variance in the estimation of the likelihood. To be sure our Monte Carlo estimators for the likelihood are trustworthy, we study the impact of Monte Carlo uncertainty in every inference we perform. Specifically, we compute the variance in the log-likelihood estimator, which varies across parameter space due to our resampling techniques in

Equation (A1) and Equation (A2). Propagating the uncertainty in the log-likelihood along independent degrees of freedom, the variance in the log-likelihood estimator $\sigma^2_{\ln\hat{\mathcal{L}}}$ in combining Equation (2), Equation (3), and Equation (A1) can be estimated as (e.g., Essick & Farr 2022)

$$\sigma^2_{\ln\hat{\mathcal{L}}}(\mathbf{\Lambda}) = \sum_{i=1}^{N_{\rm det}} \frac{\sigma^2_{\hat{\mathcal{L}}_i}(\mathbf{\Lambda})}{\hat{\mathcal{L}}_i(\mathbf{\Lambda})^2} + N^2_{\rm det}\sigma^2_{\hat{\xi}}(\mathbf{\Lambda}), \tag{A3}$$

where

$$\sigma^2_{\hat{\mathcal{L}}_i}(\mathbf{\Lambda}) = \frac{1}{N_{\rm PE}}\left[\frac{1}{N_{\rm PE}-1}\sum_{j=1}^{N_{\rm PE}}\left(\frac{\pi(\boldsymbol{\theta}_{ij}|\mathbf{\Lambda})}{p(\boldsymbol{\theta}_{ij})}\right)^2 - \hat{\mathcal{L}}_i(\mathbf{\Lambda})^2\right] \tag{A4}$$

is the Monte Carlo variance in the single event Monte Carlo integrals of Equation (A1) and

$$\sigma^2_{\hat{\xi}}(\mathbf{\Lambda}) = \frac{1}{N_{\rm draw}}\left[\frac{1}{N_{\rm draw}-1}\sum_{j=1}^{N_{\rm found}}\left(\frac{\pi(\boldsymbol{\theta}_j|\mathbf{\Lambda})}{p(\boldsymbol{\theta}_j|\mathbf{\Lambda}_{\rm draw})}\right)^2 - \hat{\xi}(\mathbf{\Lambda})^2\right] \tag{A5}$$

is the variance in the detection efficiency Monte Carlo integral of Equation (A2). When the rate-marginalized likelihood of Equation (4) is used, $\sigma^2_{\ln\hat{\mathcal{L}}}$ takes a slightly different form

$$\sigma^2_{\ln\hat{\mathcal{L}}}(\mathbf{\Lambda}) = \sum_{i=1}^{N_{\rm det}} \frac{\sigma^2_{\hat{\mathcal{L}}_i}(\mathbf{\Lambda})}{\hat{\mathcal{L}}_i(\mathbf{\Lambda})^2} + N^2_{\rm det}\frac{\sigma^2_{\hat{\xi}}(\mathbf{\Lambda})}{\hat{\xi}(\mathbf{\Lambda})^2}. \tag{A6}$$

It has been shown that GW population inference can be biased when the variance in log-likelihood estimator exceeds 1. Consequently, we adopt a threshold on $\sigma^2_{\ln\hat{\mathcal{L}}}$ of 1 to manage the bias of the posterior. Above this threshold the likelihood estimate may not be converged, and thus we ignore posterior samples with variances above this chosen threshold. Equation (A3) describes the pointwise variance in the estimation of the log-likelihood. However, for accurate sampling of the posterior we only require that the difference of log-likelihoods to be small. A small pointwise log-likelihood variance is *sufficient* for the variance of the difference of log-likelihoods to be small but not *necessary*, rendering our threshold conservative (Farr 2019; Essick & Farr 2022). Indeed, for some models, a large region of hyperparameter space is removed by this threshold, limiting the range of potential populations that can be explored; for an example, see Appendix D.3. Improvements to likelihood estimation are an active area of ongoing research (Wysocki et al. 2019; Doctor et al. 2019; Delfavero et al. 2021; Golomb & Talbot 2022; Mould et al. 2024; Hussain et al. 2024; Mancarella & Gerosa 2025).

A.2. *Sampling Techniques*

In each model, we draw samples from the posterior to study the population distributions consistent with the data and the population model. We draw samples using a variety of stochastic sampling algorithms, where the exact approach depends on the model. For most strongly modeled approaches, we use the nested sampler Dynesty (Speagle 2020) wrapper in Bilby (Ashton et al. 2019), and the GWPopulation implementation of the hierarchical likelihood (Talbot et al. 2025a).

However, our weakly modeled approaches tend to have a large number of hyperparameters and so have a high dimensional posterior. Nested sampling struggles with high dimensional distributions, so we use the Hamiltonian Monte Carlo (HMC) adaptive No-U-Turn Sampler (NUTS) implementation in NumPyro (Phan et al. 2019; Bingham et al. 2019). The NumPyro adaptive NUTS requires an autodifferentiable implementation of the likelihood, which we write in Jax (Bradbury et al. 2018). This gradient information allows the NUTS algorithm to efficiently explore high dimensional posteriors.

B. SUMMARY OF MODELS USED IN THE STRONGLY MODELED APPROACH

B.1. *Mass Model for the Full CBC Population*

Our strongly modeled FullPop-4.0 mass model is based on Power Law+Dip (Fishbach et al. 2020), Broken Power Law + Dip (Farah et al. 2022), and MultiPDB (Mali & Essick 2025). It is also similar to the Power

LAW+DIP+BREAK model (Abbott et al. 2023a), but in our FULLPOP-4.0 model we now include additional structure. We allow for the possibility of an upper mass gap, as well as separate pairing functions for NS-containing binaries and BBHs. Additionally, we include explicit peaks at low and mid-range BH masses to capture the ~ 9–$10\,M_\odot$ and ~ 30–$40\,M_\odot$ features found by the weakly modeled approaches. As discussed in Section 4, we find that the addition of a low-mass BH peak eliminates the need to explicitly model a gap between NSs and BHs masses.

Our mass model is parameterized as

$$\pi(m_1, m_2|\mathbf{\Lambda}) \propto \pi_m(m_1|\mathbf{\Lambda})\pi_m(m_2|\mathbf{\Lambda})f(m_1, m_2)\Theta(m_2 < m_1) \tag{B7}$$

where the one-dimensional mass distribution $\pi_m(m|\mathbf{\Lambda})$ is given by

$$\begin{aligned}\pi_m(m|\mathbf{\Lambda}) = \; & [1 + c_1\mathcal{N}_{[m_{\rm min},m_{\rm max}]}(m|\mu_1,\sigma_1) + c_2\mathcal{N}_{[m_{\rm min},m_{\rm max}]}(m|\mu_2,\sigma_2)]n_1(m\,|\,m_{\rm NSmax}, m_{\rm BHmin}, \eta_{\rm NSmax}, \eta_{\rm BHmin}, A) \\ & \times n_2(m\,|\,m_{\rm UMGmin}, m_{\rm UMGmax}, \eta_{\rm UMGmin}, \eta_{\rm UMGmax}, A_2)h(m\,|\,m_{\rm NSmin}, \eta_{\rm NSmin})l(m\,|\,m_{\rm BHmax}, \eta_{\rm BHmax}) \\ & \times \begin{cases} m^{\alpha_1} & \text{if } m < m_{\rm NSmax} \\ m^{\alpha_{\rm dip}} m_{\rm NSmax}^{\alpha_1-\alpha_{\rm dip}} & \text{if } m_{\rm NSmax} \le m < m_{\rm BHmin} \\ m^{\alpha_2} m_{\rm NSmax}^{\alpha_1-\alpha_{\rm dip}} m_{\rm BHmin}^{\alpha_{\rm dip}-\alpha_2} & \text{if } m \ge m_{\rm BHmin}. \end{cases}\end{aligned} \tag{B8}$$

This $\pi_m(m|\mathbf{\Lambda})$ represents a universal mass function to describe the primary and secondary mass distributions. Note that the marginal mass distribution is different from the universal mass distribution due to the pairing formalism. $\mathcal{N}_{[a,b]}(\mu,\sigma)$ represents a truncated normal distribution over $[a,b]$ with location and width parameters μ and σ. The high-pass, low-pass and notch functions are defined as follows:

$$l(m|m_{\rm BHmax}, \eta_{\rm BHmax}) = \left[1 + \left(\frac{m}{m_{\rm BHmax}}\right)^{\eta_{\rm BHmax}}\right]^{-1}, \tag{B9}$$

$$h(m|m_{\rm NSmin}, \eta_{\rm NSmin}) = 1 - l(m|m_{\rm NSmin}, \eta_{\rm NSmin}), \tag{B10}$$

$$n_1(m|m_{\rm NSmax}, m_{\rm BHmin}, \eta_{\rm NSmax}, \eta_{\rm BHmin}, A) = 1 - Al(m|m_{\rm NSmax}, \eta_{\rm NSmax})h(m|m_{\rm BHmin}, \eta_{\rm BHmin}), \tag{B11}$$

$$n_2(m|m_{\rm UMGmin}, m_{\rm UMGmax}, \eta_{\rm UMGmin}, \eta_{\rm UMGmax}, A_2) = 1 - A_2 l(m|m_{\rm UMGmin}, \eta_{\rm UMGmin})h(m|m_{\rm UMGmax}, \eta_{\rm UMGmax}). \tag{B12}$$

$$\tag{B13}$$

The pairing function

$$f(m_1, m_2|\beta_{\rm BH}, \beta_{\rm NS}) = \begin{cases} \left(\dfrac{m_2}{m_1}\right)^{\beta_1} & \text{if } m_2 < 5\,M_\odot \\ \left(\dfrac{m_2}{m_1}\right)^{\beta_2} & \text{if } m_2 > 5\,M_\odot \end{cases} \tag{B14}$$

controls how much merging binaries favor/disfavor equal masses. We allow for alternative pairing for binaries with very light secondary masses (NSBHs or BBH with the secondary component in the lower-mass gap, e.g., GW190814 Abbott et al. 2020e). We show the priors and describe the parameters of the FULLPOP-4.0 model in Table 4.

B.2. *Neutron Star Mass Models*

Following previous work (Landry & Read 2021; Abbott et al. 2023a), the mass distribution of NS-containing events is modeled as

$$\pi(m_1, m_2|\mathbf{\Lambda}) \propto \begin{cases} \pi(m_1|\mathbf{\Lambda})\,\pi(m_2|\mathbf{\Lambda}) & \text{if BNS,} \\ \mathrm{U}(3\,M_\odot, 60\,M_\odot)\,\pi(m_2|\mathbf{\Lambda}) & \text{if NSBH,} \end{cases} \tag{B15}$$

To model $\pi(m|\mathbf{\Lambda})$, we use either of the following models

Table 4. Summary of FullPop-4.0 model parameters and priors.

Category	Parameter	Unit	Description	Prior
Pairing Function	β_1	–	Spectral index below 5 $M_\odot$	U$(-2, 3)$
	β_2	–	Spectral index above 5 $M_\odot$	U$(-2, 7)$
Broken Power-Law	α_1	–	Powerlaw below $m_{\rm NS\ max}$	U$(-10, 2)$
	$\alpha_{\rm dip}$	–	Powerlaw between $m_{\rm NS\ max}$ and $m_{\rm BH\ min}$	U$(-3, 2)$
	α_2	–	Powerlaw above $m_{\rm BH\ min}$	U$(-3, 2)$
	$m_{\rm brk}$	$M_\odot$	Break point between α_1 and α_2	5
Highpass Filter	$m_{\rm NS\ min}$	$M_\odot$	Low-mass roll-off	U$(1, 1.4)$
	$\eta_{\rm min}$	–	Sharpness at $m_{\rm NS\ min}$	50
Lowpass Filter	$m_{\rm BHmax}$	$M_\odot$	High-mass roll-off	U$(60, 200)$
	$\eta_{\rm max}$	–	Sharpness at $m_{\rm BHmax}$	U$(-4, 12)$
Low-Mass Notch	$m_{\rm NS\ max}$	$M_\odot$	Lower notch edge	U$(1.4, 5)$
	$\eta_1^{\rm low}$	–	Sharpness at $m_{\rm NS\ max}$	50
	$m_{\rm BH\ min}$	$M_\odot$	Upper notch edge	U$(5, 9)$
	$\eta_1^{\rm high}$	–	Sharpness at $m_{\rm BH\ min}$	50
	A_1	–	Notch depth	0
High-Mass Notch	$m_{\rm UMGmin}$	$M_\odot$	Lower notch edge	U$(30, 90)$
	$\eta_2^{\rm low}$	–	Sharpness at $m_{\rm UMGmin}$	30
	$m_{\rm UMGmax}$	$M_\odot$	Upper notch edge	U$(60, 150)$
	$\eta_2^{\rm high}$	–	Sharpness at $m_{\rm UMGmax}$	30
	A_2	–	Depth of high-mass notch	U$(0, 1)$
Low-Mass Peak	$\mu_2^{\rm peak}$	$M_\odot$	Peak location	U$(6, 12)$
	$\sigma_2^{\rm peak}$	$M_\odot$	Peak width	U$(0, 5)$
	c_2	–	Peak height	U$(0, 500)$
High-Mass Peak	$\mu_1^{\rm peak}$	$M_\odot$	Peak location	U$(17, 50)$
	$\sigma_1^{\rm peak}$	$M_\odot$	Peak width	U$(4, 20)$
	c_1	–	Peak height	U$(0, 1000)$

Note—U(x, y) denotes a Uniform prior between x and y.

Table 5. Summary of Power and Peak NS mass model parameters.

Parameter	Unit	Description	Prior
α	–	Spectral index for the power-law in the Power NS mass distribution.	U$(-15, 5)$
$m_{\rm min}$	$M_\odot$	Minimum mass of the NS mass distribution.	U$(1.0, 1.5)$
$m_{\rm max}$	$M_\odot$	Maximum mass of the NS mass distribution.	U$(1.5, 3.0)$
μ	$M_\odot$	Location of the Gaussian peak in the Peak NS mass distribution.	U$(1.0, 3.0)$
σ	$M_\odot$	Width of the Gaussian peak in the Peak NS mass distribution.	U$(0.01, 2.00)$

Note—U(x, y) denotes a Uniform prior between x and y.

1. Power model:

$$\pi(m|\mathbf{\Lambda}) \propto \begin{cases} m^{\alpha} & \text{if } m_{\min} \leq m \leq m_{\max}, \\ 0 & \text{otherwise.} \end{cases} \tag{B16}$$

2. Peak model:

$$\pi(m|\mathbf{\Lambda}) \propto \begin{cases} \exp\left[-\dfrac{(m-\mu)^2}{2\sigma^2}\right] & \text{if } m_{\min} \leq m \leq m_{\max}, \\ 0 & \text{otherwise.} \end{cases} \tag{B17}$$

See Table 5 for a description of the parameters used in the model and the corresponding prior ranges.

B.3. *Binary Black Hole Mass Models*

Broken Power Law + 2 Peaks: The fiducial BBH mass model is a mixture between a broken power law and two left-truncated Gaussian peaks, with low mass tapering applied to the full distribution. The broken power law is given by

$$p_{\rm BP}(m_1|\alpha_1, \alpha_2, m_{\rm break}, m_{1,\rm low}, m_{\rm high}) = \frac{1}{N} \begin{cases} \left(\dfrac{m_1}{m_{\rm break}}\right)^{-\alpha_1} & m_{1,\rm low} \leq m_1 < m_{\rm break} \\ \left(\dfrac{m_1}{m_{\rm break}}\right)^{-\alpha_2} & m_{\rm break} \leq m_1 < m_{\rm high}, \end{cases} \tag{B18}$$

where α_1 and α_2 are the power law indices, the transition between the low-mass and high-mass power law occurs at $m_{\rm break}$, and the normalization constant is

$$N = \frac{m_{\rm high}(m_{\rm high}/m_{\rm break})^{-\alpha_2} - m_{\rm break}}{1-\alpha_2} + \frac{m_{\rm high}(m_{\rm high}/m_{\rm break})^{-\alpha_1} - m_{\rm break}}{1-\alpha_1}. \tag{B19}$$

The full mixture distribution $\pi(m_1|\mathbf{\Lambda})$ is

$$\begin{aligned} \pi(m_1|\mathbf{\Lambda}) \propto \Big[& \lambda_0 p_{\rm BP}(m_1|\alpha_1, \alpha_2, m_{\rm break}, m_{1,\rm low}, m_{\rm high}) + \lambda_1 \mathcal{N}_{lt}(m_1|\mu_1, \sigma_1, \text{low} = m_{1,\rm low}) \\ & + (1-\lambda_0-\lambda_1)\mathcal{N}_{lt}(m_1|\mu_2, \sigma_2, \text{low} = m_{1,\rm low}) \Big] S(m_1|m_{1,\rm low}, \delta_{m,1}), \end{aligned} \tag{B20}$$

where $\mathcal{N}_{lt}$ is a left-truncated normal distribution. The Planck tapering function S ensures a smooth turn-on of the distribution in the range $(m_{1,\rm low}, m_{1,\rm low} + \delta_{m,1}]$ and is given by

$$S(m|m_{\rm low}, \delta_m) = \begin{cases} 0 & m < m_{\rm low}, \\ [1 + f(m - m_{\rm low}, \delta_m)]^{-1} & m_{\rm low} \leq m < m_{\rm low} + \delta_m, \\ 1 & m_{\rm low} + \delta_m \leq m, \end{cases} \tag{B21}$$

with

$$f(m', \delta_m) = \exp\left(\frac{\delta_m}{m'} + \frac{\delta_m}{m' - \delta_m}\right).$$

We model the mass ratio as a power law with index β_q and low-mass tapering applied to secondary mass m_2, conditioned on primary mass m_1,

$$p_{\rm PL}(q|m_1, \beta_q, m_{2,\rm low}, \delta_{m,2}) \propto q^{\beta_q} S(m_1 q|m_{2,\rm low}, \delta_{m,2}), \tag{B22}$$

with the same tapering function defined above. We numerically normalize the mass ratio distribution as a function of primary mass, and interpolate the normalization to arbitrary primary masses with a log-uniform grid across primary

Table 6. Summary of Broken Power Law + 2 Peaks model parameters and priors.

Parameter	Description	Prior
α_1	Spectral index of 1st primary mass power law	$\mathrm{U}(-4, 12)$
α_2	Spectral index of 2nd primary mass power law	$\mathrm{U}(-4, 12)$
m_{break}	Power law break location	$\mathrm{U}(20, 50)$
μ_1	Location of the first peak	$\mathrm{U}(5, 20)$
σ_1	Width of the first peak	$\mathrm{U}(0, 10)$
μ_2	Location of the second peak	$\mathrm{U}(25, 60)$
σ_2	Width of the second peak	$\mathrm{U}(0, 10)$
$m_{1,\mathrm{low}}$	Lower edge of taper function	see (B23)
$\delta_{\mathrm{m},1}$	Mass range of low mass tapering	$\mathrm{U}(0, 10)$
λ_0, λ_1	Mixing fractions between power law and peaks	$\mathrm{Dir}(\alpha = (1, 1, 1))$
m_{high}	Maximum mass for distribution, which is pinned to $m_{\mathrm{high}} = 300\,M_\odot$ by default	$\delta(m_{\mathrm{high}} - 300)$
β_q	Spectral index of mass ratio power law	$\mathrm{U}(-2, 7)$
$m_{2,\mathrm{low}}$	Lower edge of taper function in m_2	see (B23)
$\delta_{\mathrm{m},2}$	Mass range of low mass tapering in m_2	$\mathrm{U}(0, 10)$

Note—The priors for the Extended Broken Power Law + 2 Peaks model are the same except for the β_q parameters, which assume a prior of $\mathrm{U}(-10, 10)$.

mass. We define $m_2 \leq m_1$, and therefore must enforce $m_{2,\mathrm{low}} \leq m_{1,\mathrm{low}}$. We use a prior which is uniform in the two dimensional triangular space satisfying the inequality and between 3 and 10 $M_\odot$. Specifically, this defines the priors

$$\begin{aligned} \pi(m_{1,\mathrm{low}}) &= \frac{2}{(\mathrm{max} - \mathrm{min})^2}(m_{1,\mathrm{low}} - \mathrm{min}), \\ \pi(m_{2,\mathrm{low}}|m_{1,\mathrm{low}}) &= \frac{1}{m_{1,\mathrm{low}} - \mathrm{min}}, \end{aligned} \tag{B23}$$

where $\mathrm{min} = 3\,M_\odot, \mathrm{max} = 10\,M_\odot$. The priors used in this model can be found in Table 6.

Extended Broken Power Law + 2 Peaks: This model incorporates correlations between the primary mass and mass ratio, by allowing each primary mass mixture component in the Broken Power Law + 2 Peaks to be associated with a separate power law mass ratio model. In terms of the functions defined above, the model is expressed as

$$\begin{aligned} \pi(m_1, q|\mathbf{\Lambda}) \propto \Big[& \lambda_0 p_{\mathrm{BP}}(m_1|\alpha_1, \alpha_2, m_{\mathrm{break}}, m_{1,\mathrm{low}}, m_{\mathrm{high}}) p_{\mathrm{PL}}(q|m_1, \beta_q^{\mathrm{BP}}, m_{2,\mathrm{low}}^{\mathrm{BP}}, \delta_{m,2}^{\mathrm{BP}}) \\ & + \lambda_1 \mathcal{N}_{lt}(m_1|\mu_1, \sigma_1, \mathrm{low} = m_{1,\mathrm{low}}) p_{\mathrm{PL}}(q|m_1, \beta_q^{\mathrm{peak1}}, m_{2,\mathrm{low}}^{\mathrm{peak1}}, \delta_{m,2}^{\mathrm{peak1}}) \\ & + (1 - \lambda_0 - \lambda_1) \mathcal{N}_{lt}(m_1|\mu_2, \sigma_2, \mathrm{low} = m_{1,\mathrm{low}}) p_{\mathrm{PL}}(q|m_1, \beta_q^{\mathrm{peak2}}, m_{2,\mathrm{low}}^{\mathrm{peak2}}, \delta_{m,2}^{\mathrm{peak2}}) \Big] S(m_1|m_{1,\mathrm{low}}, \delta_{m,1}), \end{aligned} \tag{B24}$$

where the superscripts (BP, peak1, and peak2) denote the different mass ratio hyperparameters (β_q, $m_{2,\mathrm{low}}$, $\delta_{m,2}$) for the broken power law, first Gaussian, and second Gaussian components. The priors on these hyperparameters are the same as those used in the Broken Power Law + 2 Peaks model, listed in Table 6, except for the power law index parameters β_q, which assume $\mathrm{U}(-10, 10)$.

Table 7. Summary of Power Law Redshift model parameter and prior.

Parameter	Description	Prior
κ	Power-law index on comoving merger rate evolution	$\mathrm{U}(-10, 10)$

B.4. Power Law Redshift *Model*

We model redshift evolution by the comoving merger rate density (Equation 5). In particular, we use the model

$$\pi(z|\kappa) \propto \frac{1}{1+z}\frac{\mathrm{d}V_c}{\mathrm{d}z}(1+z)^{\kappa} \tag{B25}$$

where the prefactor converts from a rate density in comoving volume and source frame time to detector frame time and redshift. In other words, the comoving rate density scales as $\mathcal{R} \propto (1+z)^{\kappa}$. We use a prior on κ as specified in Table 7.

B.5. *Spin Models*

BBH spins can be parameterized in several ways which are useful for GW data analysis. Here, we use dimensionless spin magnitudes for the BHs χ_1 and χ_2, and cosine tilt angles $\cos\theta_1$ and $\cos\theta_2$, as well as the effective spin parameters $\chi_{\rm eff}$ (Racine 2008; Ajith et al. 2011; Damour 2001) and $\chi_{\rm p}$ (Schmidt et al. 2011, 2012, 2015). The effective inspiral spin $\chi_{\rm eff}$ used because it is typically the most precisely measured BBH spin parameter, due to its lower-order appearance post-Newtonian expansions (Arun et al. 2009). Other parametrizations of spin precession exist (e.g., Fairhurst et al. 2020; Gerosa et al. 2021; Thomas et al. 2021) but for reasons of convention, we only work with the typical parameterization given in Equation 16 of Abac et al. (2025a). When modeling the effective spins, we use the analytic per-event parameter-estimation prior on $\chi_{\rm eff}$, $\chi_{\rm p}$, and q (Iwaya et al. 2025) in the calculation of the hierarchical likelihood (Equation 2). For all models presented in this work, we assume that azimuthal angles ϕ_i are distributed uniformly between 0 and 2π—the same as their parameter estimation prior—due to their typically uninformative individual-event posteriors. See Table 3 of Abac et al. (2025a) for more information about spin parameters.

We report results with spin vectors defined at the reference frequencies used for inference on each signal (Abac et al. 2025b). The effective spin $\chi_{\rm eff}$ is approximately conserved throughout the inspiral (Racine 2008; Gerosa et al. 2015), so its population distribution should be unaffected by the choice of reference frequency. While the effective precessing spin $\chi_{\rm p}$ and the spin angles are dependent on reference frequency, they are comparatively weakly constrained by the data. Additionally, their parameter estimation priors are invariant under time evolution, meaning measurements are robust against different reference frequencies. At the current number of events and because of model dependence while measuring tilt distributions, we do not expect significant differences between the inferred tilt or $\chi_{\rm p}$ distributions at different reference frequencies (Mould & Gerosa 2022). Our approach is consistent with past LVK analyses, where evolved spins have not been used (Abbott et al. 2019a, 2021a, 2023a).

B.5.1. *Component Spin Models*

Gaussian Component Spins: We model the spin magnitudes (χ_i) as a truncated Gaussian distribution between 0 and 1, assuming they are identically and independently distributed:

$$\pi(\chi_i|\mu_\chi, \sigma_\chi) = \mathcal{N}_{[0,1]}(\chi_1|\mu_\chi, \sigma_\chi)\mathcal{N}_{[0,1]}(\chi_2|\mu_\chi, \sigma_\chi)\,. \tag{B26}$$

This choice attempts to rectify shortcomings of the non-singular Beta distribution used to model the spin magnitudes in previous work (Abbott et al. 2019a, 2021a, 2023a); see Equation (B29) below. The non-singular Beta distribution is forced to go to $\pi(\chi) = 0$ at $\chi = 0, 1$, which crucially does not allow for measurements of contributions to the population at $\chi = 0$ or $\chi = 1$, even though non-trivial contributions may exist (Callister et al. 2022; Galaudage et al. 2021; Hussain et al. 2024). Allowing the Beta distribution to be singular would add limited additional model flexibility, only adding the option for $\pi(\chi) = \infty$ at the $\chi = 0, 1$, but not anywhere in between 0 and ∞. Thus, we opt for the truncated Gaussian model, which can take on a continuous range of values at χ's boundaries.

Table 8. Summary of GAUSSIAN COMPONENT SPINS model parameters for spin magnitudes (Equation B26) and tilt angles (Equation B27, B28).

Parameter	Description	Prior
μ_χ	Location of the χ distribution	U(0, 1)
σ_χ	Width of the χ distribution	U(0.005, 1)
μ_t	Location of the Gaussian component of the $\cos\theta$ distribution	U(−1, 1)
σ_t	Width of the Gaussian component of the $\cos\theta$ distribution	U(0.01, 4)
ζ	Fraction in the Gaussian component of the $\cos\theta$ distribution	U(0, 1)
$t_{\rm min}$	Minimum of the $\cos\theta$ distribution	U(−1, 1)

NOTE—U stands for a uniform prior.

We model the distribution of the cosine spin tilt angle ($\cos\theta_i$) as a mixture between a Gaussian distribution truncated on −1 to 1 and an isotropic distribution, assuming they are identically but *not* independently distributed:

$$\pi(\cos\theta_i|\mu_t,\sigma_t,\zeta) = \zeta\,\mathcal{N}_{[-1,1]}(\cos\theta_1|\mu_t,\sigma_t)\mathcal{N}_{[-1,1]}(\cos\theta_2|\mu_t,\sigma_t) + \frac{1-\zeta}{4}\,. \tag{B27}$$

We here allow for the location of the Gaussian sub-population to vary, following Vitale et al. (2022), rather than fixing it at $\mu_t = 1$ as was done in previous work (Abbott et al. 2019a, 2021a, 2023a). Priors on the GAUSSIAN COMPONENT SPINS hyperparameters are given in Table 8.

MINIMUM TILT MODEL: For the spin tilts, we also use a model in which $t_{\rm min}$ (where $t \equiv \cos\theta$), a hard lower truncation of $\pi(\cos\theta)$, is informed by the data. In this case the population distribution becomes

$$\pi(\cos\theta_i|\mu_t,\sigma_t,\zeta,t_{\rm min}) = \begin{cases} \zeta\,\mathcal{N}_{[t_{\rm min},1]}(\cos\theta_1|\mu_t,\sigma_t)\mathcal{N}_{[t_{\rm min},1]}(\cos\theta_2|\mu_t,\sigma_t) + \dfrac{1-\zeta}{(1-t_{\rm min})^2} & \cos\theta_i > t_{\rm min}\ , \\ 0 & \cos\theta_i \leq t_{\rm min}\ . \end{cases} \tag{B28}$$

Equation (B27) is the same as Equation (B28) with $t_{\rm min} = -1$. The prior on $t_{\rm min}$ and other hyperparameters are given in Table 8.

BETA DISTRIBUTION SPIN MAGNITUDE: In Appendix D.2, we compare the GAUSSIAN COMPONENT SPINS model to alternatives. For the spin magnitudes, these are the Constrained and Unconstrained Beta Distributions, which both follow the form:

$$\pi(\chi_i|\alpha,\beta) \propto {\chi_i}^{\alpha-1}(1-\chi_i)^{\beta-1}\ . \tag{B29}$$

The Constrained Beta was the DEFAULT model for GWTC-3.0 (Abbott et al. 2023a) and requires the shape parameters $\alpha, \beta > 1$, which forces $\pi(\chi_i = 0, 1) = 0$ (Wysocki et al. 2019; Abbott et al. 2021a). The Unconstrained Beta relaxes this constraint on the the shape parameters. If $0 < \alpha, \beta < 1$, the Beta distribution becomes singular, meaning $\pi(\chi_i = 0, 1) = \infty$. Following Abbott et al. (2023a), we sample the mean and standard deviation of the Beta distribution, rather than α and β, and use the same priors for μ_χ and σ_χ as given in Table 8 for the GAUSSIAN COMPONENT SPINS model. For the Constrained Beta case, we additionally impose the cut $\alpha, \beta > 1$. The relationship between $\{\mu_\chi,\ \sigma_\chi\}$ and $\{\alpha,\ \beta\}$ is given in Equation 5 of Abbott et al. (2019a).

B.5.2. *Identical versus Non-identically Distributed Spin Magnitudes and Tilts*

In Appendix D.2, we investigate whether or not the primary and secondary spins are identically distributed, assuming the GAUSSIAN COMPONENT SPINS model. *Identical* distribution means that the two distributions share the same set of hyperparameters, while *non-identical* means that each is described by different hyperparameters. Spin magnitudes are also *independently* distributed, meaning that $\pi(\chi_1, \chi_2)$ is separable in terms of χ_1 and χ_2. The acronym IID means they are independently and identically distributed, while IND means independently and non-identically distributed:

$$\chi_i \text{ IID} \implies \pi(\chi_{1,2}|\mu_\chi,\sigma_\chi) = \mathcal{N}_{[0,1]}(\chi_1|\mu_\chi,\sigma_\chi)\,\mathcal{N}_{[0,1]}(\chi_2|\mu_\chi,\sigma_\chi)\ , \tag{B30}$$

$$\chi_i \text{ IND} \implies \pi(\chi_{1,2}|\mu_{\chi,1},\sigma_{\chi,1},\mu_{\chi,2},\sigma_{\chi,2}) = \mathcal{N}_{[0,1]}(\chi_1|\mu_{\chi,1},\sigma_{\chi,1})\,\mathcal{N}_{[0,1]}(\chi_2|\mu_{\chi,2},\sigma_{\chi,2})\ . \tag{B31}$$

Table 9. Summary of Gaussian Effective Spins (Equation B36) and Skew-normal Effective Spin (Equation B37) spin parameters.

Parameter	Description	Prior G	Prior SN
$\mu_{\rm eff}$	Location of the $\chi_{\rm eff}$ distribution	U(−1, 1)	U(−1, 1)
$\sigma_{\rm eff}$	Width of the $\chi_{\rm eff}$ distribution	U(0.05, 1)	LU(0.01, 4)
μ_p	Location of the $\chi_{\rm p}$ distribution	U(0.05, 1)	U(0.01, 1)
σ_p	Width of the $\chi_{\rm p}$ distribution	U(0.07, 1)	LU(0.01, 1)
ρ	Degree of correlation between $\chi_{\rm eff}$ and $\chi_{\rm p}$	U(−0.75, 0.75)	N/A
ϵ	Skew of the $\chi_{\rm eff}$ distribution	N/A	U(−1, 1)

Note—The first column of priors (**G**) gives those used for the Gaussian Effective Spins results, the second column (**SN**) is Skew-normal Effective Spin. U stands for a uniform prior; LU for log-uniform.

The tilt angles, on the other hand, are *non-independently* distributed, meaning that $\pi(\cos\theta_1, \cos\theta_2)$ is not separable. The acronynm NID means they are non-independently but identically distributed, while NND means non-independently and non-identically distributed:

$$\cos\theta_i \text{ NID} \implies \pi(\cos\theta_{1,2}|\zeta,\mu_t,\sigma_t) = \zeta\,\mathcal{N}_{[-1,1]}(\cos\theta_1|\mu_t,\sigma_t)\,\mathcal{N}_{[-1,1]}(\cos\theta_2|\mu_t,\sigma_t) + \frac{1-\zeta}{4}\ , \tag{B32}$$

$$\cos\theta_i \text{ NND} \implies \pi(\cos\theta_{1,2}|\zeta,\mu_{t,1},\sigma_{t,1},\mu_{t,2},\sigma_{t,2}) = \zeta\,\mathcal{N}_{[-1,1]}(\cos\theta_1|\mu_{t,1},\sigma_{t,1})\,\mathcal{N}_{[-1,1]}(\cos\theta_2|\mu_{t,2},\sigma_{t,2}) + \frac{1-\zeta}{4}\ . \tag{B33}$$

For the non-identically distributed cases, each component's hyperparameters have the same priors as those listed in Table 8 for the identically distributed case. Table 12 gives Bayes factors between different combinations of IID/IND spin magnitudes and NID/NND spin tilts; see associated dicsussion in Appendix D.2.

B.5.3. *Effective Spin Models*

Gaussian Effective Spins: We here assume that the distribution of $\chi_{\rm eff}$ and $\chi_{\rm p}$ across the BBH population is a bivariate truncated Gaussian (Miller et al. 2020; Roulet & Zaldarriaga 2019) characterized by the location and width of the $\chi_{\rm eff}$ and $\chi_{\rm p}$ distributions, and the covariance between them:

$$\pi(\chi_{\rm eff},\chi_{\rm p}|\mu,\mathbf{\Sigma}) \propto \mathcal{N}(\chi_{\rm eff},\chi_{\rm p}|\mu,\mathbf{\Sigma})\ , \tag{B34}$$

where $\mu = (\mu_{\rm eff},\mu_{\rm p})$ and

$$\mathbf{\Sigma} = \begin{pmatrix} \sigma_{\rm eff}^2 & \rho\,\sigma_{\rm eff}\,\sigma_{\rm p} \\ \rho\,\sigma_{\rm eff}\,\sigma_{\rm p} & \sigma_{\rm p}^2 \end{pmatrix}\ . \tag{B35}$$

This expands to:

$$\pi(\chi_{\rm eff},\chi_{\rm p}|\mu,\mathbf{\Sigma}) \propto \exp\left[-\frac{1}{2(1-\rho^2)}\left(\frac{(\chi_{\rm eff}-\mu_{\rm eff})^2}{\sigma_{\rm eff}^2} - \frac{2\rho(\chi_{\rm eff}-\mu_{\rm eff})(\chi_{\rm p}-\mu_{\rm p})}{\sigma_{\rm eff}\sigma_{\rm p}} + \frac{(\chi_{\rm p}-\mu_{\rm p})^2}{\sigma_{\rm p}^2}\right)\right]\ . \tag{B36}$$

The bivariate Gaussian is truncated over the range $\chi_{\rm eff}\in[-1,1]$, $\chi_{\rm p}\in[0,1]$ and is normalized numerically. Priors on the hyperparameters are given in Table 9.

Skew-normal Effective Spin: To account for observational evidence that the $\chi_{\rm eff}$ distribution is not symmetric (Callister et al. 2021b; Adamcewicz & Thrane 2022; Banagiri et al. 2025), we additionally model $\chi_{\rm eff}$ as a skewed, truncated Gaussian distribution:

$$\pi(\chi_{\rm eff}|\mu_{\rm eff},\sigma_{\rm eff},\epsilon) \propto \begin{cases} (1+\epsilon)\,\mathcal{N}_{[-1,1]}(\chi_{\rm eff}|\mu_{\rm eff},\sigma_{\rm eff}(1+\epsilon)) & \chi_{\rm eff}\leq 0\ , \\ (1-\epsilon)\,\mathcal{N}_{[-1,1]}(\chi_{\rm eff}|\mu_{\rm eff},\sigma_{\rm eff}(1-\epsilon)) & \chi_{\rm eff}\geq 0\ . \end{cases} \tag{B37}$$

Table 10. Summary of parameters exclusive to the COPULA, LINEAR, and SPLINE correlated models, along with their priors.

Parameter	Description	Prior
$\kappa_{x,y}$	Level of correlation between parameters x and y inferred with a copula	$\mathrm{U}(-20,20)$
$\mu_{\mathrm{eff}\|q}$	$q=1$ intercept in linear $\mu_{\mathrm{eff}}(q)$	$\mathrm{U}(-1,1)$
$\delta\mu_{\mathrm{eff}\|q}$	Gradient in linear $\mu_{\mathrm{eff}}(q)$	$\mathrm{U}(-2,2)$
$\ln\sigma_{\mathrm{eff}\|q}$	$q=1$ intercept in linear $\ln\sigma_{\mathrm{eff}}(z)$	$\mathrm{U}(-5,0)$
$\delta\ln\sigma_{\mathrm{eff}\|q}$	Gradient in linear $\ln\sigma_{\mathrm{eff}}(q)$	$\mathrm{U}(-12,4)$
$\mu_{\mathrm{eff}\|z}$	$z=0$ intercept in linear $\mu_{\mathrm{eff}}(z)$	$\mathrm{U}(-1,1)$
$\delta\mu_{\mathrm{eff}\|z}$	Gradient in linear $\mu_{\mathrm{eff}}(z)$	$\mathrm{U}(-1,1)$
$\ln\sigma_{\mathrm{eff}\|z}$	$z=0$ intercept in linear $\ln\sigma_{\mathrm{eff}}(z)$	$\mathrm{U}(-5,0)$
$\delta\ln\sigma_{\mathrm{eff}\|z}$	Gradient in linear $\ln\sigma_{\mathrm{eff}}(z)$	$\mathrm{U}(-3,5)$
$\mu^i_{\mathrm{eff}\|q}$	i^{th} node in spline $\mu_{\mathrm{eff}}(q)$	$\mathrm{U}(-1,1)$
$\ln\sigma^i_{\mathrm{eff}\|q}$	i^{th} node in spline $\ln\sigma_{\mathrm{eff}}(q)$	$\mathrm{U}(-5,0)$
$\mu^i_{\mathrm{eff}\|z}$	i^{th} node in spline $\mu_{\mathrm{eff}}(z)$	$\mathrm{U}(-1,1)$
$\ln\sigma^i_{\mathrm{eff}\|z}$	i^{th} node in spline $\ln\sigma_{\mathrm{eff}}(z)$	$\mathrm{U}(-5,0)$

NOTE— There are four variations of the COPULA model in which $(x,y)=(q,\chi_{\mathrm{eff}})$, $(m_1,\chi_{\mathrm{eff}})$, (z,χ_{eff}) and finally, (m_1,z). Within the spline models, we use four nodes i.

This model includes a parameter ϵ which describes the skew of the distribution, such that $\epsilon>0$ characterizes a distribution with more support for $\chi_{\mathrm{eff}}<\mu_{\mathrm{eff}}$, while $\epsilon<0$ has more support when $\chi_{\mathrm{eff}}>\mu_{\mathrm{eff}}$. The SKEW-NORMAL EFFECTIVE SPIN distribution reduces to a standard, symmetric Gaussian in the case that $\epsilon=0$. We here model χ_{p} as a truncated normal distribution and infer its location and width. Priors on the hyperparameters are given in Table 9.

B.6. *Copula Correlation Models*

Copulas allow for a variable correlation between two parameters x and y while keeping the marginal distributions for the respective parameters fixed. This is done using the fact that the cumulative distribution function (CDF) of any randomly distributed variable is itself a uniformly distributed variable between 0 and 1, in order to model the CDFs of x and y with a correlated two-dimensional uniform distribution. This correlated two-dimensional uniform distribution, known as a copula density function, is dependent on a (hyper)parameter $\kappa_{x,y}$, which determines the level (or strength) of the correlation. A coordinate transformation (which is determined by the chosen marginal distributions for x and y) can then be applied to provide a correlated two-dimensional model in the desired coordinates. Conveniently, the Jacobian for this transformation turns out to be the product of the chosen marginal distributions, meaning the two-dimensional population model can be written as

$$\pi_{xy}(x,y|\mathbf{\Lambda},\kappa_{x,y}) = \pi_c\left(u(x|\mathbf{\Lambda}),v(y|\mathbf{\Lambda})|\kappa_{x,y}\right)\pi_x(x|\mathbf{\Lambda})\pi_y(y|\mathbf{\Lambda}). \tag{B38}$$

Here, $\mathbf{\Lambda}$ are the set of hyperparameters governing the marginal distributions π_x and π_y, while π_c denotes the chosen copula density function. Finally,

$$u(x|\mathbf{\Lambda}) = \int_{x_{\mathrm{min}}}^{x} \mathrm{d}x'\ \pi_x(x'|\mathbf{\Lambda}), \tag{B39}$$

and

$$v(y|\mathbf{\Lambda}) = \int_{y_{\mathrm{min}}}^{y} \mathrm{d}y'\ \pi_y(y'|\mathbf{\Lambda}), \tag{B40}$$

which we refer to as u and v for the sake of brevity, are the CDFs of x and y respectively.

The COPULA models all assume the same marginal distributions and copula density functions, but differ in which pairs of parameters they correlate (i.e., which parameters u and v are functions of in the above notation). Namely,

the mass distribution is assumed to follow the default Broken Power Law + 2 Peaks model, redshift follows the default power-law redshift model, and $\chi_{\rm eff}$ and $\chi_{\rm p}$ are Gaussian distributed, but uncorrelated (i.e., they follow the Gaussian Effective Spins model with $\rho = 0$). We assume a Frank copula density function

$$\pi_c\left(u, v|\kappa_{x,y}\right) = \frac{-\kappa_{x,y} e^{-\kappa_{x,y}(u+v)}\left(e^{-\kappa_{x,y}} - 1\right)}{\left(e^{-\kappa_{x,y}} - e^{-\kappa_{x,y}u} - e^{-\kappa_{x,y}v} + e^{-\kappa_{x,y}(u+v)}\right)^2}, \tag{B41}$$

where $\kappa_{x,y} \in (-\infty, \infty)$. Positive values of $\kappa_{x,y}$ imply a correlation, while negative values of $\kappa_{x,y}$ imply an anti-correlation. The Frank copula density function is not defined at $\kappa_{x,y} = 0$, but becomes uncorrelated as $\kappa_{x,y} \to 0$ from above and below, so we assume $\pi_c\left(u, v|\kappa_{x,y} = 0\right) = 1$ (technically, making this a piece-wise function). The Frank copula density function is chosen as it allows for positive and negative correlations that are symmetric about $\kappa_{x,y} = 0$, and gives rise to correlated distributions that we believe appear physically reasonable (Adamcewicz & Thrane 2022; Adamcewicz et al. 2023). The prior on $\kappa_{x,y}$ (which is the same for all copula model variations), is given in Table 10.

Copulas are advantageous as they allow for a potential correlation to be quantified by a single variable $\kappa_{x,y}$, that otherwise has no influence on the distribution of the population. However, relative to the linear and spline models explored below, they suffer from a lack of flexibility when it comes to covariance. There are a limited number of copula density functions available, all of which introduce a correlation in a unique, but rigid way (e.g, Adamcewicz & Thrane 2022; Adamcewicz et al. 2023). Furthermore, known two-dimensional copulas depend on a single parameter $\kappa_{x,y}$, and cannot infer, for example, a separable broadening and correlation with the mean of a distribution simultaneously. As a result, using copulas to probe for more complex structure in two-dimensions (assuming fixed marginal distributions), requires model comparison between a number of different copula density functions (e.g., Adamcewicz et al. 2023).

B.7. *Linear Correlation Models*

The $(q, \chi_{\rm eff})$ and $(z, \chi_{\rm eff})$ Linear models begin by assuming that $\chi_{\rm eff}$ is Gaussian distributed for any given value of mass ratio q and redshift z respectively. As the two models are otherwise identical, for the remainder of this Section, we substitute x for q and z. This parameter-dependent Gaussian distribution for $\chi_{\rm eff}$, $\pi(\chi_{\rm eff}|x)$, is truncated at unphysical values of $|\chi_{\rm eff}| \geq 1$. From here, the mean and (natural log) width of the $\chi_{\rm eff}$ distribution are allowed to evolve with the chosen variable x linearly:

$$\mu_{\rm eff}(x) = \mu_{{\rm eff}|x} + \delta\mu_{{\rm eff}|x} x, \tag{B42}$$

and

$$\ln \sigma_{\rm eff}(x) = \ln \sigma_{{\rm eff}|x} + \delta \ln \sigma_{{\rm eff}|x_0} x. \tag{B43}$$

Here, $\mu_{{\rm eff}|x}$, $\delta\mu_{{\rm eff}|x}$, $\ln \sigma_{{\rm eff}|x}$, and $\delta \ln \sigma_{{\rm eff}|x_0}$ are all hyperparameters fit to the data, where $\delta\mu_{{\rm eff}|x}$ and $\delta \ln \sigma_{{\rm eff}|x_0}$ quantify the strength of the correlation between x and the mean and width of the $\chi_{\rm eff}$ distribution respectively. The priors for these parameters are given in Table 10. Meanwhile, masses follow the default Broken Power Law + 2 Peaks model and redshift is distributed according to the default power-law redshift model.

These models (in the style of those presented in Safarzadeh et al. 2020; Callister et al. 2021b; Biscoveanu et al. 2022a), therefore allow us to infer separable trends in the mean and width of the $\chi_{\rm eff}$ distribution with another parameter. The inferred correlations should be interpreted with care as, unlike the copula models defined in Section B.6, the correlation hyperparameters here will also affect the shape of the marginal $\chi_{\rm eff}$ distribution. As such, it can be difficult to be certain whether an inferred correlation is entirely due to a trend between $\chi_{\rm eff}$ and another parameter, or is in part due to a better fit to the marginal $\chi_{\rm eff}$ distribution (Adamcewicz & Thrane 2022).

B.8. *Spline Correlation Models*

The Spline models (Heinzel et al. 2024) are similar to the Linear models defined in Section B.7. The key difference is that the Spline models more flexibly model the parameter-dependent mean $\mu_{\rm eff}(x)$ and (natural log) width $\ln \sigma_{\rm eff}(x)$ of the $\chi_{\rm eff}$ distribution as cubic splines dependent on x. Each spline has four nodes $\mu^i_{\rm eff}$ and $\ln \sigma^i_{{\rm eff}|x}$ that are placed uniformly in $\log_{10} x$, and are inferred from the data. The priors for these nodes are given in Table 10.

C. SUMMARY OF MODELS USED IN THE WEAKLY MODELED APPROACH

We describe weakly modeled approaches—B-Spline, BGP, AR, FM—used in this work. Results in the main text make use of B-Spline and BGP models. BGP was chosen since it was the only model capable of modeling the full mass spectrum, while B-Spline was chosen for most BBH analyses for its ease of description and flexibility, as it

simultaneously models all parameters with B-Splines. See Appendix D.4 for comparison of these weakly modeled approaches.

C.1. B-Spline *models*

B-Spline: The B-Spline model (Edelman et al. 2023) simultaneously fits all population parameters with Basis-splines (B-splines). A k^{th} order B-spline of variable x consists of a linear combination of n basis functions $\{B_{k,n}(x)\}$, each of which are degree $k-1$ piecewise polynomials joined at a set of m locations called knots $\{x_m\}$. The number of basis functions n defined across any range of x values is determined solely by the order of the spline k and the total number of knots m as $n = k + m$. This forms a basis that spans the space of possible interpolants between the knots $\{x_m\}$. Given a vector of coefficients $\boldsymbol{\alpha}$, any function f can then be approximated as

$$\tilde{f}(x) = \sum_{i=1}^{n} B_{k,i}(x)\alpha_i. \tag{C44}$$

In the case of this work, f is a probability distribution $p(\theta|\boldsymbol{\alpha}_\theta)$ for a population-level parameter θ, where the B-spline coefficients $\boldsymbol{\alpha}_\theta$ are the hyperparameters of the distribution that are inferred during parameter estimation.

Given an adequate number of knots, a B-spline is highly flexible and therefore capable of identifying sharp features that could be present in the data. This does, however, naturally make the B-Spline model prone to overfitting. To combat this, we include a smoothing prior π_{BS} that penalizes large differences between neighboring coefficients. The prior is defined as

$$p(\boldsymbol{\alpha}|\tau) = \exp\left(-\frac{1}{2}\tau\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{D}_r^{\mathrm{T}}\boldsymbol{D}_r\boldsymbol{\alpha}\right), \tag{C45}$$

where D_r is the r-order difference matrix with shape $(n-r)\times n$ and τ is a scalar that controls the level of smoothing. Ideally, τ is also inferred during parameter estimation. We found that τ consistently railed against prior boundaries, which meant that the limits imposed on the likelihood uncertainty was the main driver of smoothness. We therefore fix τ to a reasonable value (roughly between 5−10) for each population distribution. With a sufficient number of bases, typically $n \sim 30$−40, this penalty prior will prevent the B-spline from overfitting the data while the large number of bases will provide enough flexibility to fit sharp features.

The mass and spin distributions are modeled as B-splines, and the redshift distribution is modeled as a power law modulated by a B-spline (Edelman et al. 2023), with all components inferred simultaneously. One of the weakly modeled approaches used in previous analyses (Abbott et al. 2023a) was the Powerlaw + Spline model. This model assumed that the primary BBH mass followed a power-law distribution with moderate deviations controlled by a cubic spline. The B-Spline model does not assume an underlying shape for the primary mass distribution, instead allowing for full model flexibility.

The B-Spline model infers the separable components of the mass distributions, $p(m_1)$ and $p(q)$, wherein the primary mass is defined over the range 3−300 $M_\odot$ and the mass ratio is defined over the range 0.03−1. Unlike the Broken Power Law + 2 Peaks model, a minimum secondary mass is not enforced during parameter estimation. To provide a more direct comparison to the strongly modeled approach, the B-Spline mass distributions shown in Section 6.1 are not the separable components $p(m_1)$ or $p(q)$ but instead the marginal distributions conditioned on $m_2 > 3\,M_\odot$, that is, $p(q|m_2 > 3\,M_\odot) = \int p(q)p(m_1)\Theta(m_1 q - 3)\mathrm{d}m_1$. We include the separable distributions along with the conditional marginal distributions in the mass ratio plot in Figure 20 to illustrate how this assumption affects the shape of the mass ratio distribution.

Isolated Peak: This model is defined by two subpopulations (Godfrey et al. 2023). One assumes a primary mass distribution described by a log-Gaussian peak while the other infers the mass distribution with a B-spline. Mass ratio, spin magnitude, and spin tilt distributions are inferred separately for each subpopulation, also using B-splines. The redshift distribution is the same for each subpopulation and is inferred with a power law modulated by a B-spline.

C.2. Binned Gaussian Process *(BGP) Model*

In previous population analyses (Abbott et al. 2023a), we used the BGP model to study the joint distribution of primary and secondary masses. Here, we extend it to model the joint distributions of mass, spin, and redshift. The BGP approach models the rate of BBH, BNS and NSBH mergers as a piecewise constant function over a set of fixed bins across the one-, two-, or three-dimensional joint space. The comoving merger rate density in each bin is a

hyperparameter of the model. In addition, the BGP model couples the logarithm of the rate density in each bin with a Gaussian process covariance, assuming an exponential quadratic kernel (also known as a radial basis function (RBF) kernel). The exponential quadratic kernel has a hyper-hyperparameter length scale λ for each parameter in the joint space.

The covariance between two bins is proportional to the exponential of the negative squared distance between the bin centers, in units of the length scales along each parameter. In addition, there is one more hyper-hyperparameter σ, which acts as an overall multiplicative scaling of the covariance matrix (Ray et al. 2023a). The BGP approach directly infers the rate density in each bin, as well as the hyper-hyperparameters of the covariance kernel. The prior on the parameter σ is a half-normal with width 1, and the length scale priors were tuned to the mean and variation in the set of distances between the bin centers.

We used 22 bins spaced uniformly in log-mass for the mass BGP models, and 15 bins spaced uniformly between $\chi_{\rm eff} \in [-0.7, 0.7]$ for the effective spin BGP models.

The BGP approach has the advantage of being able to model nontrivial correlations in the population of compact binaries. However, it assumes an arbitrary binning scheme, which reduces the resolution of the constraints and leads to unphysical discontinuities at the boundary between bins. Furthermore, measurements of the rate density in some bins may be driven by the *a priori* assumption of a Gaussian covariance, and not the data. This can also cause smoothing near sharp features e.g., near the $m_1 = m_2$ boundary.

C.3. Autoregressive Process *(AR) Model*

The AR model (Callister & Farr 2024) is a highly flexible model for one dimensional marginal merger rate densities. The Monte Carlo integrals for estimating the likelihood in Equation (2) involve a large set of parameter estimation samples [Equation (A1)] and found injections for the selection efficiency [Equation (A2)]. The merger rate density at each sample is its own hyperparameter, which is directly inferred from the data. Without any further assumptions, such a model is severely underconstrained and will converge on the maximum likelihood functional distribution (Payne & Thrane 2023). In order to *a priori* favor smoother distributions, this finite list of samples is ordered along the dimension of interest like, for example, their primary mass. The marginal merger rate density is then assumed to be an autoregressive Gaussian random walk in log-space. There are two hyper-hyperparameters, $\sigma_{\rm AR}$ and $\tau_{\rm AR}$, which control the scale of variability and the autocorrelation length along the Gaussian random walk respectively. These hyper-hyperparameters have half-normal and log-normal priors respectively, tuned to the scale of the data: see Appendix B in Callister & Farr (2024) for details. These hyper-hyperparameters are jointly inferred along with the logarithmic merger rate density at each sample.

The AR model is particularly well-suited to distributions with sharp features and nontrivial evolution, complementing the BGP approach. However, the AR model used here cannot model correlations in the population. Furthermore, as in the BGP approach, the inferred merger rate may extrapolate based off nearby constraints in regions of little information. When evaluating uncertainties in the AR model, one should consider the impact of the AR smoothing prior. In particular, lower bounds on the merger rate in some regions may not be an accurate lower bound on the true distribution.

Like the B-Spline model, the AR model does not enforce a minimum secondary mass through a mass-dependent mass ratio.

C.4. Flexible Mixtures *(FM) Model*

The FM model (Tiwari 2021) is a flexible mixture model framework for modeling the BBH population. FM models the population as a Gaussian mixture model in chirp mass and aligned spins, and a power law in mass ratio and redshift (Tiwari 2022). Furthermore, FM is able to model correlations between parameters in the population, and additionally has a variable model dimension, sampling the trans-dimensional posterior using a reversible-jump MCMC method. For the analyses presented here, FM uses 11 mixture components.

D. RESULT VALIDATION STUDIES

In the following sub-appendices, we describe various methods used to select and validate the data, models, and approaches presented in the main text. In Appendix D.1 and Appendix D.2 we discuss the process for selecting the fiducial mass and spin models, respectively, under the strongly modeled approach and present results from other models which were under consideration. Appendix D.3 describes how the specific convergence cut biases the inferred

Table 11. Bayes factor comparison of selected models from the mass model comparison study.

Model	Abbreviation	$\log_{10}\mathcal{B}$
Broken Power Law + 2 Peaks ★	BP2P	0
Broken Power Law + 1 Peak	BP1P	−0.06
Broken Power Law + 3 Peaks	BP3P	−0.34
Power Law + Peak	PP	−2.43

effective spin distribution. In Appendix D.4, we compare various types of methods for the weakly modeled approach, which were presented in Appendix C. Appendix D.5 gives compact object merger rates when sub-threshold triggers are considered, i.e., using a lower FAR threshold. Finally, Appendix D.6 provides supplementary results for population-level correlations between masses, spins, and redshifts.

D.1. *Model Comparison Study: Mass*

The fiducial mass model is a strongly modeled approach that is intended to provide a minimal but accurate description of the data with a parametrization that is more readily interpretable than the weakly modeled approaches. Table 11 shows the Bayes factors between the models that performed the best in our study and our fiducial model Broken Power Law + 2 Peaks. We also compare to our previous fiducial model from Abbott et al. (2023a), the Power Law + Peak model. In Figure 15, we see that Power Law + Peak infers the peak component at $35M_\odot$, as in Abbott et al. (2023a), but a broken power law plus 1 peak (BP1P) model infers the peak component at $10M_\odot$ and the power law break at $35M_\odot$. In Table 11, we see that all models that include a broken power law are strongly favored over the Power Law + Peak model. Among models with a broken power law, the Bayes factors do not show strong support for one model over another. While Bayes factors are an important model comparison statistic to consider, they can be influenced by prior assumptions. In the case of the models in Table 11 where Bayes factors are not decisive, we ultimately chose as our fiducial model the model that was a minimal extension of the Power Law + Peak model and probed features shared among all of the weakly modeled approaches shown in Figure 20. The Broken Power Law + 2 Peaks model best fit these criteria, though with some nuances that we discuss below.

In GWTC-3.0, the Power Law + Peak model identified an overdensity in the merger rate at $m_1 = 35^{+1.7}_{-2.9}\,M_\odot$ relative to a global power law. Evidence for this feature first appeared in Abbott et al. (2019a) and was strengthened in Abbott et al. (2021a) and Abbott et al. (2023a). As mentioned in Section 6.1, this feature at $35M_\odot$ may be an overdensity relative to an underlying mass continuum or could mark the onset of a decline in the merger rate. In addition to the Bayes factor comparison shown in Table 11, this ambiguity is also present in the Broken Power Law + 2 Peaks posterior, which includes two modes that correspond to two different morphologies, shown in Figure 16. The dominant mode is correlated with a narrower ($\sigma_2 = 3.1^{+2.7}_{-2.0}\,M_\odot$) peak at $\mu_2 = 33^{+2.8}_{-2.1}\,M_\odot$ and accounts for 71% of the posterior volume, while the subdominant mode is correlated with a wider ($\sigma_2 = 7.3^{+2.3}_{-2.5}\,M_\odot$) peak at $\mu_2 = 29^{+2.9}_{-3.6}\,M_\odot$ and accounts for 29% of the posterior volume. This latter mode has a morphology nearly identical to the inferred BP1P distribution (see Figure 16). The inflection point of the Gaussian, $\mu_2 + \sigma_2 = 37^{+3.7}_{-3.0}\,M_\odot$, is consistent between both modes, which indicates that the right half of the Gaussian must fall below the power law to match the declining rate in this region (the gray vertical shaded region in Figure 16).

The power law indices, α_1 and α_2 are not correlated with the two modes present in the BP2P posterior, though the measured slope of the mass distribution between $18-19\,M_\odot$ is not equivalent to α_1 in the subdominant mode. In this mode, the slope (calculated by finite difference) is $0.72^{+1.4}_{-1.7}$, which includes more support for negative values compared to the inferred power law index, $\alpha_1 = 1.7^{+1.4}_{-1.9}$. We infer the power law break location at $m_{\rm break} = 36^{+12}_{-13}\,M_\odot$. The BP1P model, which does not include a second Gaussian component, measures the break location much more precisely at $m_{\rm break} = 34.1^{+3.8}_{-3.3}\,M_\odot$. The uncertainty in the BP2P power law break is due to the presence of the second Gaussian component, which likely obscures the break most of the time. Because the peak is narrower in the dominant BP2P mode, one might expect the break location to be better constrained in this mode; however, this is not the case. The break location does not appear to be correlated with either mode.

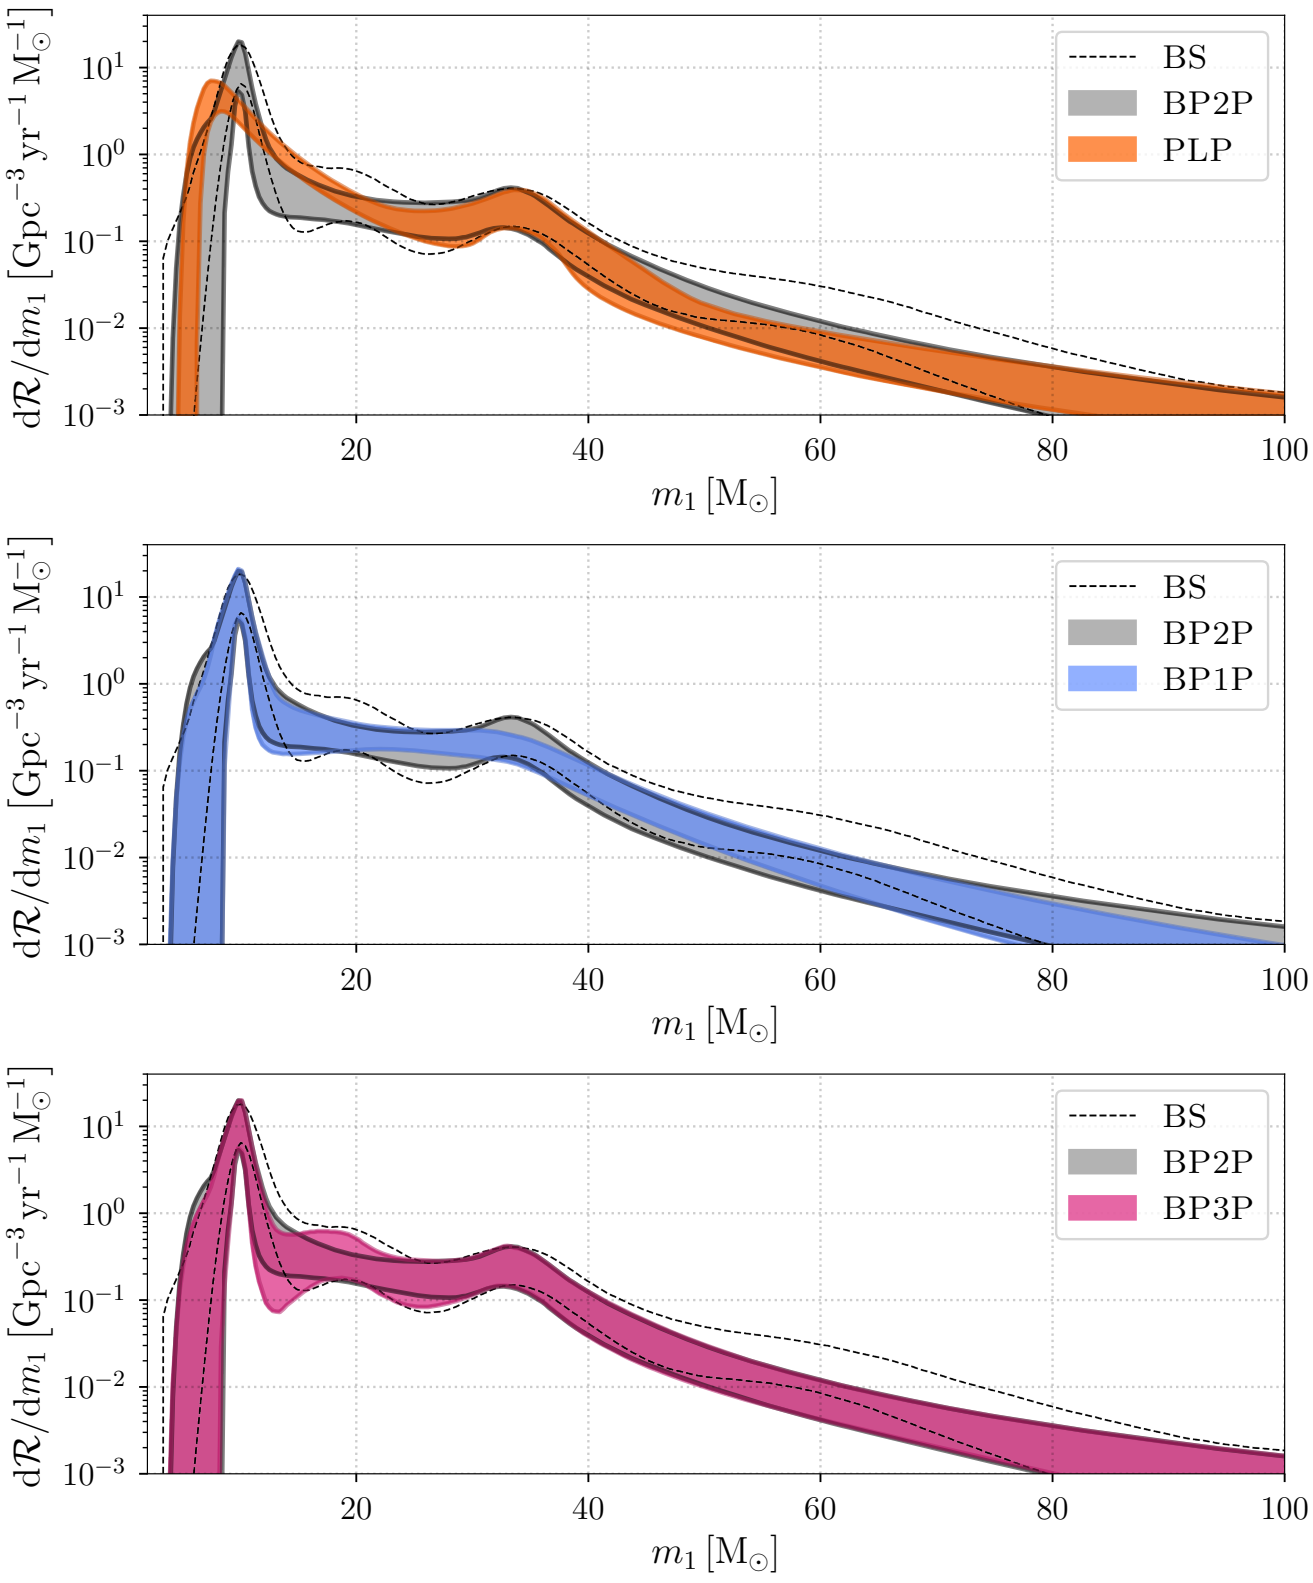

Figure 15. Inferred mass distributions of various strongly modeled approaches compared to the fiducial mass model Broken Power Law + 2 Peaks (gray shaded) and the B-Spline model (black dashed). The fiducial model from GWTC-3.0, Power Law + Peak, is shown in orange, a broken power law with a single peak in blue, and a broken power law with three peaks in pink.

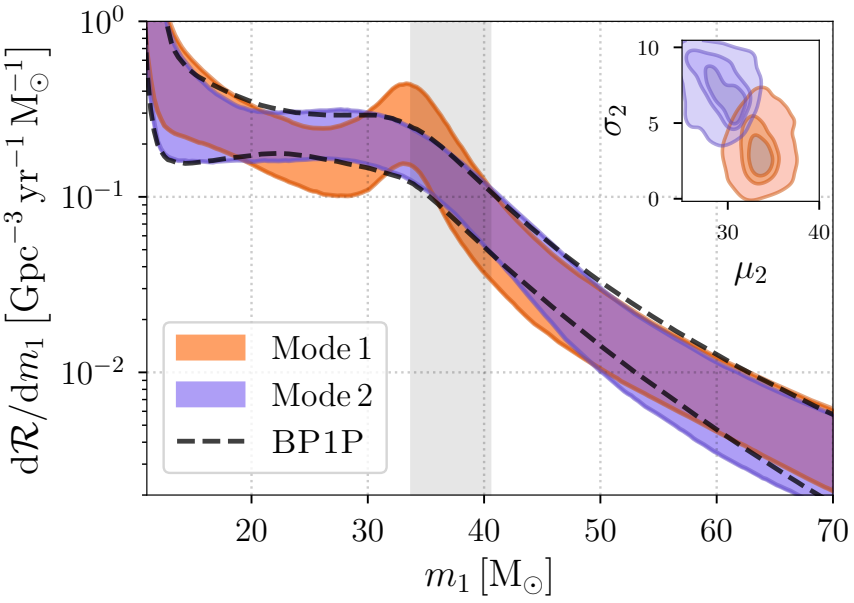

Figure 16. Differential merger rate as a function of primary mass (evaluated at $z = 0.2$) of the two different modes recovered by the Broken Power Law + 2 Peaks model. The orange shaded region shows the 90% credible interval for the dominant mode (71% of posterior), reflecting a distinct peak at 35 $M_\odot$, and the purple shaded region shows the 90% credible interval for the subdominant mode (29% of posterior), reflecting a broken power law morphology without a distinct 35 $M_\odot$ peak. The black dashed lines show the 90% credible bounds of the Broken Power Law + 1 Peak model for comparison. The inset figure shows the joint posterior of the peak mean μ_2 and width σ_2 for each mode. The vertical grey shaded region indicates the 90% credible interval of the sum $\mu_2 + \sigma_2$, which is consistent between both modes.

Table 12. Comparison of different parametric models for spin magnitudes χ_i and tilt angles $\cos\theta_i$.

χ_i model		$\cos\theta_i$ model		$\log_{10}\mathcal{B}$
Truncated Gaussian ★	IID	Isotropic + Truncated Gaussian	NID	0.0
Constrained Beta	IID	Isotropic + Aligned Gaussian	NID	−0.66
Constrained Beta	IID	Isotropic + Truncated Gaussian	NID	−0.64
Unconstrained Beta	IID	Isotropic + Truncated Gaussian	NID	−0.98
Truncated Gaussian	IID	Isotropic + Aligned Gaussian	NID	−0.08
Truncated Gaussian	IND	Isotropic + Truncated Gaussian	NID	−1.09
Truncated Gaussian	IID	Isotropic + Truncated Gaussian	NND	−0.17
Truncated Gaussian	IND	Isotropic + Truncated Gaussian	NND	−0.91

Note—The first row is the Default model for GWTC-4.0 (Gaussian Component Spins), while the second is the Default that was used for GWTC-3.0 and earlier (Abbott et al. 2019a, 2021a, 2023a; Talbot & Thrane 2017). The third through final are other models explored. All log Bayes factors ($\log_{10}\mathcal{B}$) are with respect to the first row.

D.2. *Model Comparison Study: Spin Magnitudes and Tilt Angles*

We consider two additional spin magnitude models beyond the Gaussian Component Spins model from Section 6.3.1: the Constrained Beta and Unconstrained Beta models, both defined by Equation (B29). We also consider

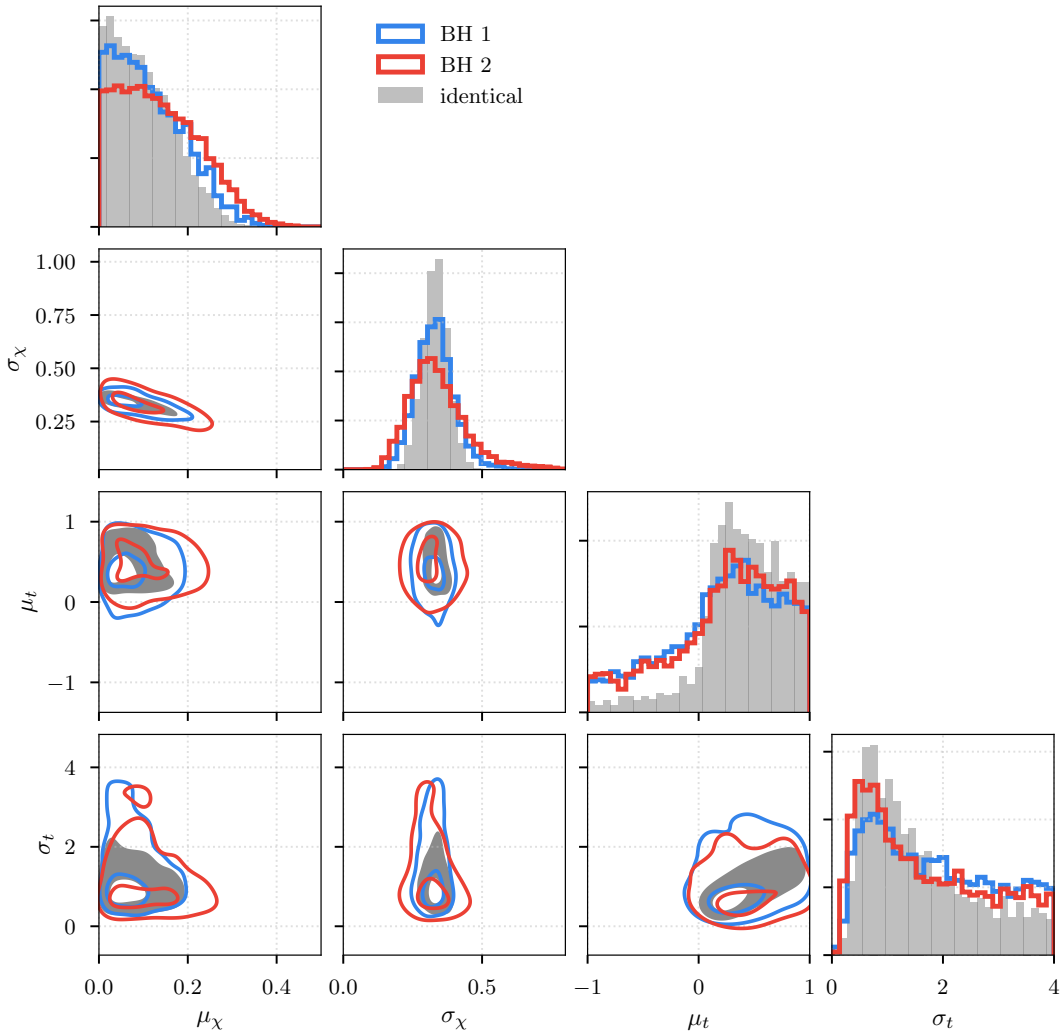

Figure 17. Distribution of Gaussian Component Spins hyperparameters for the primary BH (blue) and secondary BH (red) when neither magnitudes nor tilts are assumed identical (last row of Table 12), compared to when they are both assumed identical (gray, first row of Table 12). The contours of the two dimensional distributions mark the 50th and 90th percentiles.

one additional spin tilt model, the Isotropic + Aligned Gaussian, which fixes $\mu_t = 1$ in the Gaussian Component Spins model and served as the Default in GWTC-3.0 (Talbot & Thrane 2017; Abbott et al. 2023a). The first five rows of Table 12 give the log Bayes factors between these models. Within the Gaussian Component Spins model, we also test whether the primary and secondary spins are identically and independently distributed (see Appendix B.5.2), with log Bayes factors given in the bottom three rows of Table 12. The Gaussian Component Spins model with IID spin magnitudes and NID spin tilts performs the best, and is adopted as the default component spin model in the main text.

Figure 17 shows posteriors for the magnitude and tilt hyperparameters assuming that χ_i and $\cos\theta_i$ are both identically distributed (gray) versus both non-identically distributed (blue and red). The primary and secondary spins have consistent hyperparameter distributions. Assuming identical distribution makes the hyperparameters more-precisely constrained—yielding tighter 90% bands on the resultant population distributions—but does not affect the overall shape of the distributions. The only difference of note is that χ_2 has less support for a small μ_χ than χ_1, potentially hinting at more highly spinning secondary BHs on a population level.

D.3. *Dependence of the Effective Spin Distribution on the Likelihood Variance Cut*

We next investigate the effect of the likelihood variance threshold (see Appendix A.1) on our strongly modeled effective spin results. To ensure that posteriors generated using Monte Carlo estimation are trustworthy, i.e., the likelihood estimate is converged, we enforce that all hyperparameter samples yield a log-likelihood variance ${\sigma^2}_{\ln \hat{\mathcal{L}}} < 1$ for all analyses presented in the main text (Talbot & Golomb 2023). In GWTC-3.0, however, a different quantity was used to assess the Monte Carlo uncertainty: the effective number of independent samples, $N_{\rm eff}$ (Farr 2019). It was there imposed that $N_{\rm eff}$ be greater than $4N_{\rm events}$ for the sensitivity injections used in Equation (3) and greater than 10 for every event in the catalog of CBCs. This $N_{\rm eff}$ cut is generally less stringent than the approach taken for GWTC-4.0. The $\chi_{\rm p}$ distribution is sensitive to which method is used to cut out samples.

For the most direct comparison to GWTC-3.0, we here use the Gaussian Effective Spins model. Figure 18 shows the posteriors for the Gaussian Effective Spins parameters where the ${\sigma^2}_{\ln \hat{\mathcal{L}}}$ (purple) versus $N_{\rm eff}$ (green) cuts are done on the GWTC-4.0 results. The LVK GWTC-3.0 results (with the $N_{\rm eff}$ cut) are shown in comparison (black). Figure 19 shows the resultant marginal and joint $\chi_{\rm eff}$–$\chi_{\rm p}$ population distributions. There are three main differences between the results with the two cuts on GWTC-4.0 data:

1. The ${\sigma^2}_{\ln \hat{\mathcal{L}}}$ cut yields a $\mu_{\rm p}$ posterior which is constrained away from zero, while the the $N_{\rm eff}$ cut does not.
2. The $N_{\rm eff}$ cut allows for wider $\chi_{\rm p}$ distributions (larger $\sigma_{\rm p}$) than the ${\sigma^2}_{\ln \hat{\mathcal{L}}}$ cut.
3. Under the $N_{\rm eff}$ cut, $\chi_{\rm eff}$ and $\chi_{\rm p}$ are preferentially positively correlated, while under the ${\sigma^2}_{\ln \hat{\mathcal{L}}}$ cut we remain agnostic, with preference for small-to-zero correlation.

Thus, the claims that the $\chi_{\rm p}$ distribution does not peak at zero, and that the data prefer $\chi_{\rm eff}$ and $\chi_{\rm p}$ being uncorrelated at the population level, are not necessarily astrophysical in origin. Rather, they are driven by regions of the parameter-space that our events and sensitivity injections allow us reliably probe with Monte Carlo likelihood estimators.

The particular sensitivity of the $\chi_{\rm p}$ distribution to these cuts can be at-least partly attributed to the use of a uniform and isotropic spin prior in GWTC-4.0 parameter estimation, which induces a $\chi_{\rm p}$ prior that goes to 0 at $\chi_{\rm p} = 0, 1$. Thus, Equation (2) is prone to yield a small $N_{\rm eff}$ and/or large ${\sigma^2}_{\ln \hat{\mathcal{L}}}$ when $\chi_{\rm p}$ is near-minimal or near-maximal. Using alternative spin priors in parameter estimation could aid in getting more effective samples at small $\chi_{\rm p}$, improving our ability to probe the presence or absence of negligibly precessing BHs at the population level. The Monte Carlo uncertainty from working with a finite number of samples to estimate the likelihood will only grow more problematic as the number of observed CBCs increases (Talbot & Golomb 2023); development of mitigation techniques is an area of active research (e.g., Gerosa et al. 2020; Rinaldi & Del Pozzo 2021; Talbot & Thrane 2022; Callister et al. 2022; Talbot & Golomb 2023; Leyde et al. 2024; Hussain et al. 2024; Callister et al. 2024; Lorenzo-Medina & Dent 2025; Mancarella & Gerosa 2025).

D.4. *Comparison of Weakly Modeled Approach Results*

Mass distributions: We supplement the strongly modeled approach with various weakly modeled approaches. Figure 20 (top panel) shows the inferred primary mass distribution using the B-Spline, AR, Flexible Mixtures, and BGP models compared to the fiducial Broken Power Law + 2 Peaks model. All models agree within their 90% credible regions, though the Flexible Mixtures model is the only model that does not exhibit a prominent peak at $\sim 10 M_\odot$. The Flexible Mixtures model does not directly infer the primary mass, but instead it models the chirp mass as a Gaussian mixture model and the mass ratio with a power law. The primary mass distribution shown in Figure 20 is then derived from these two distributions. The discrepancy between this model and the other weakly modeled approaches could be due to model misspecification, i.e., assuming the mass ratio distribution follows a power law for all primary masses. The fact that the discrepancy exists primarily in the $\sim 10 M_\odot$ region supports the Isolated Peak result in Section 6.2, which suggests the $\sim 10 M_\odot$ peak disfavors equal mass mergers (i.e., is inconsistent with a power law in mass ratio) compared to higher mass BBHs. The mass ratio distributions inferred by the B-Spline, AR, and Flexible Mixtures is shown Figure 20 (bottom panel). The 90% credible regions are consistent between each model. The B-Spline model infers a peak away from $q = 1$, which is not present in the AR or Flexible Mixtures results. This may be due to the B-Spline model's greater flexibility compared to the AR and Flexible Mixtures models, as it infers all parameters simultaneously with B-Splines. The AR model assumes an auto-regressive process only in primary mass and mass ratio, while other parameters are inferred with the fiducial models listed in Table 1. As noted in Appendix C.1, the figure includes the marginal distributions conditioned on $m_2 > 3 M_\odot$ from the B-Spline and AR models in addition

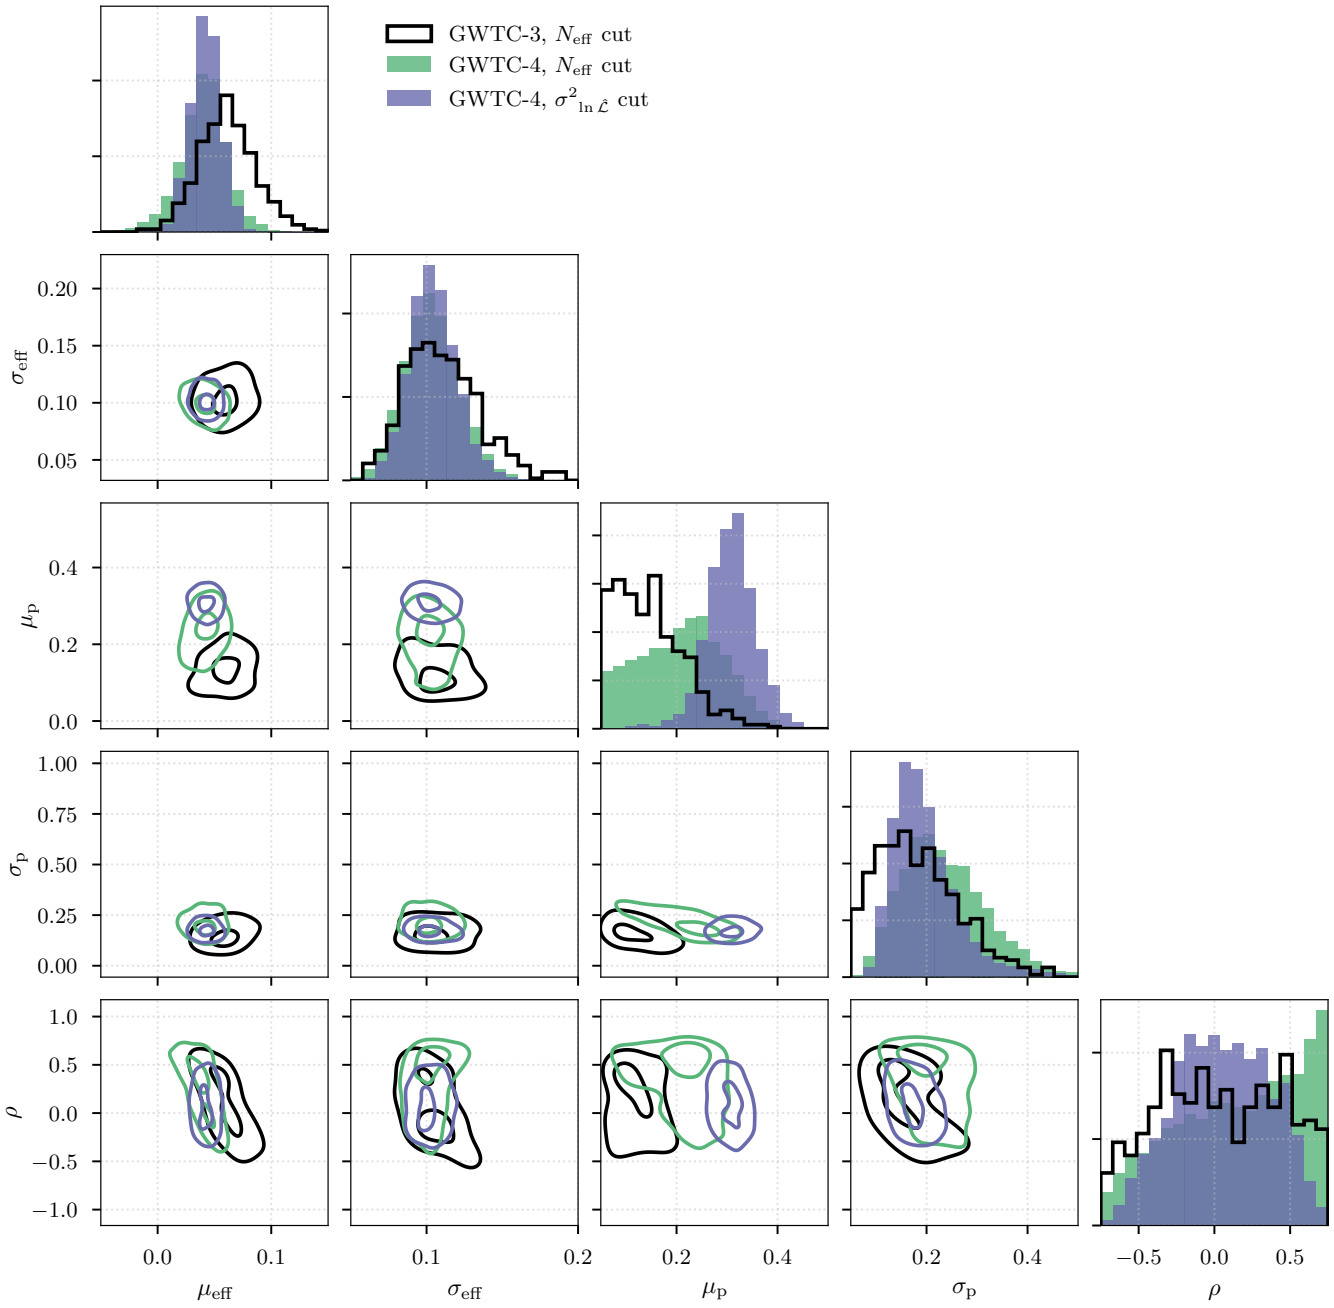

Figure 18. Posterior for the Gaussian Effective Spins model hyperparameters, as defined in Table 9, for GWTC-4.0 under two methods of cutting out samples that may lead to an unconverged likelihood. Posteriors excluding samples with a substantially large log-likelihood variance ($\sigma^2_{\ln \hat{\mathcal{L}}}$ cut) are shown in purple; those excluding samples with a substantially small number of effective samples ($N_{\rm eff}$ cut) are in green. In black are the GWTC-3.0 results with the $N_{\rm eff}$ cut for comparison. The $\chi_{\rm eff}$ hyperparameters are not affected by the cuts, while the $\chi_{\rm p}$ and joint-distribution hyperparameters are. The contours of the two dimensional distributions mark the 50th and 90th percentiles.

to the fully separable components (dashed lines). The two distributions are essentially identical above $q \sim 0.4$, below which the marginal distribution falls off sharply. This implies that the behavior of the strongly modeled approach below $q \sim 0.4$ is primarily due to the prior assumption of a minimum mass cutoff, and not due to information inferred about the shape of the distribution in that regime.

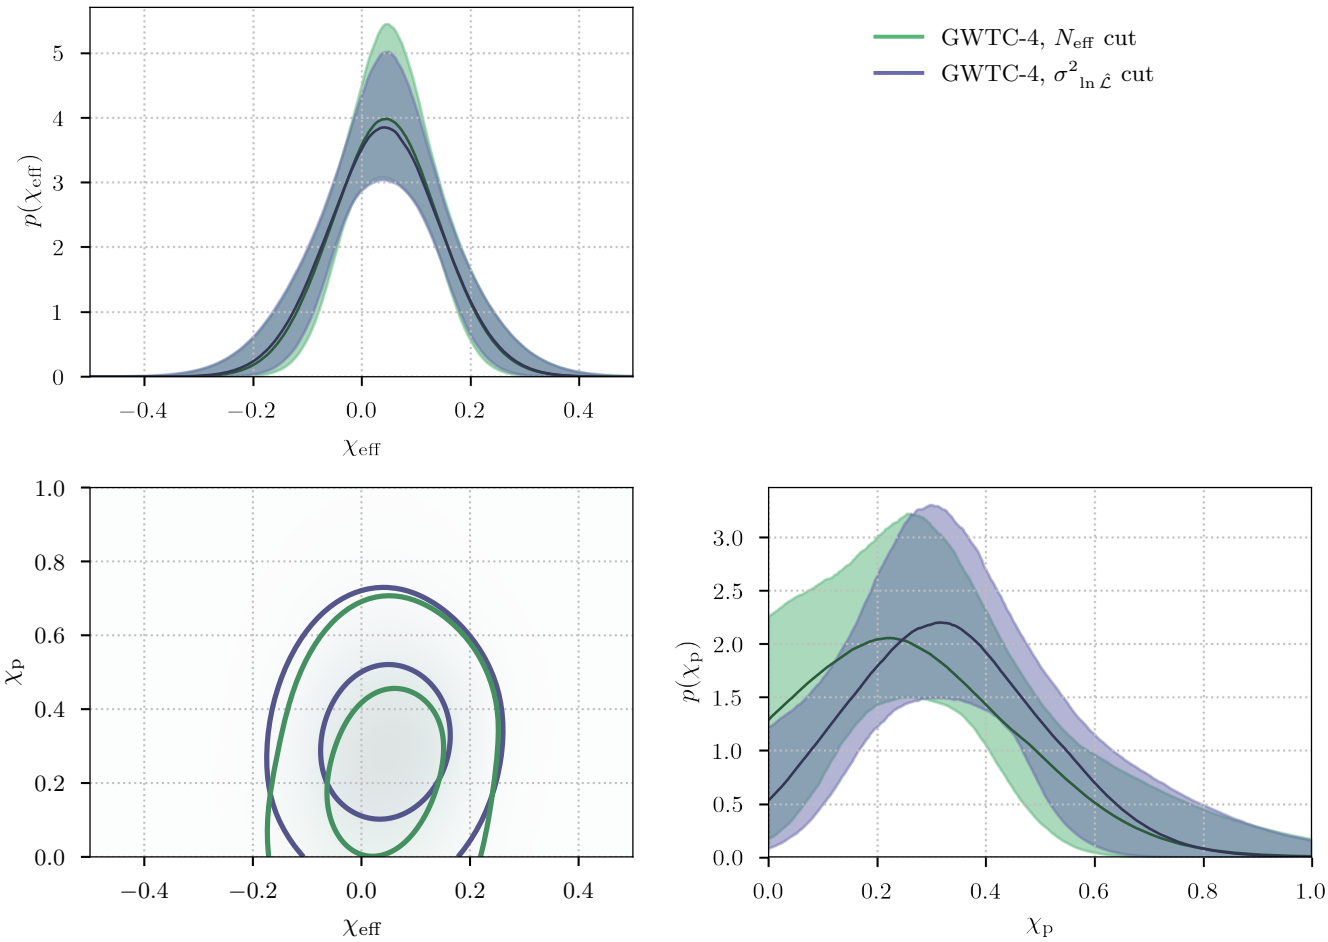

Figure 19. Marginal and joint $\chi_{\rm eff}$ and $\chi_{\rm p}$ distributions under the Gaussian Effective Spins model using two methods of cutting out hyperparameter samples that may lead to an unconverged likelihood (see Figure 18). The marginal distributions show the median (solid line) and 90% credible interval on the probability density for each population distribution (shaded). The joint distribution is the PPD with contours marking the 50th and 90th quantiles.

Effective spin distributions: In addition to the Skew-normal Effective Spin model and various correlation models presented for $\chi_{\rm eff}$ in the main text (Figures 9 and 12), we use the weakly modeled approach for a consistency check and find the distributions shown in the top panel of Figure 21. We plot the $\chi_{\rm eff}$ distribution directly inferred with a binned Gaussian process (BGP; blue) as well as that reconstructed from the spin magnitude and tilt B-Spline model results (green), and compare them to the Skew-normal Effective Spin (red-orange), Gaussian Effective Spins (purple), and $(q, \chi_{\rm eff})$ Spline (maroon) results from the main text. The two approaches paint the same qualitative picture: the $\chi_{\rm eff}$ distribution peaks at small values. However, there is variation between the weakly modeled approach and strongly modeled approach; the weakly modeled distributions are wider than the strongly modeled distributions. In Figure 21, we additionally plot posteriors distributions on statistics derived from each $\chi_{\rm eff}$ distribution (c.f., Table 3), with results summarized as follows:

- *First percentile of the* $\chi_{\rm eff}$ *distribution* ($\chi_{\rm eff,1\%}$): The three strongly modeled distributions find consistent first-percentiles, around ~ 0.2, while the weakly modeled find that the $\chi_{\rm eff}$ distribution extends to lower values. In the case of the BGP, the first-percentile distribution rails against the minimum allowed value of -0.7.

- *Fraction of BBHs with negative* $\chi_{\rm eff}$: All models yield consistent posteriors with one another and find that the fraction of BBHs with negative $\chi_{\rm eff}$ is greater than zero and less than ~ 0.8. The BGP posterior on this fraction is the widest, and the B-Spline posterior peaks at a higher values than the near-identical strongly modeled posteriors.

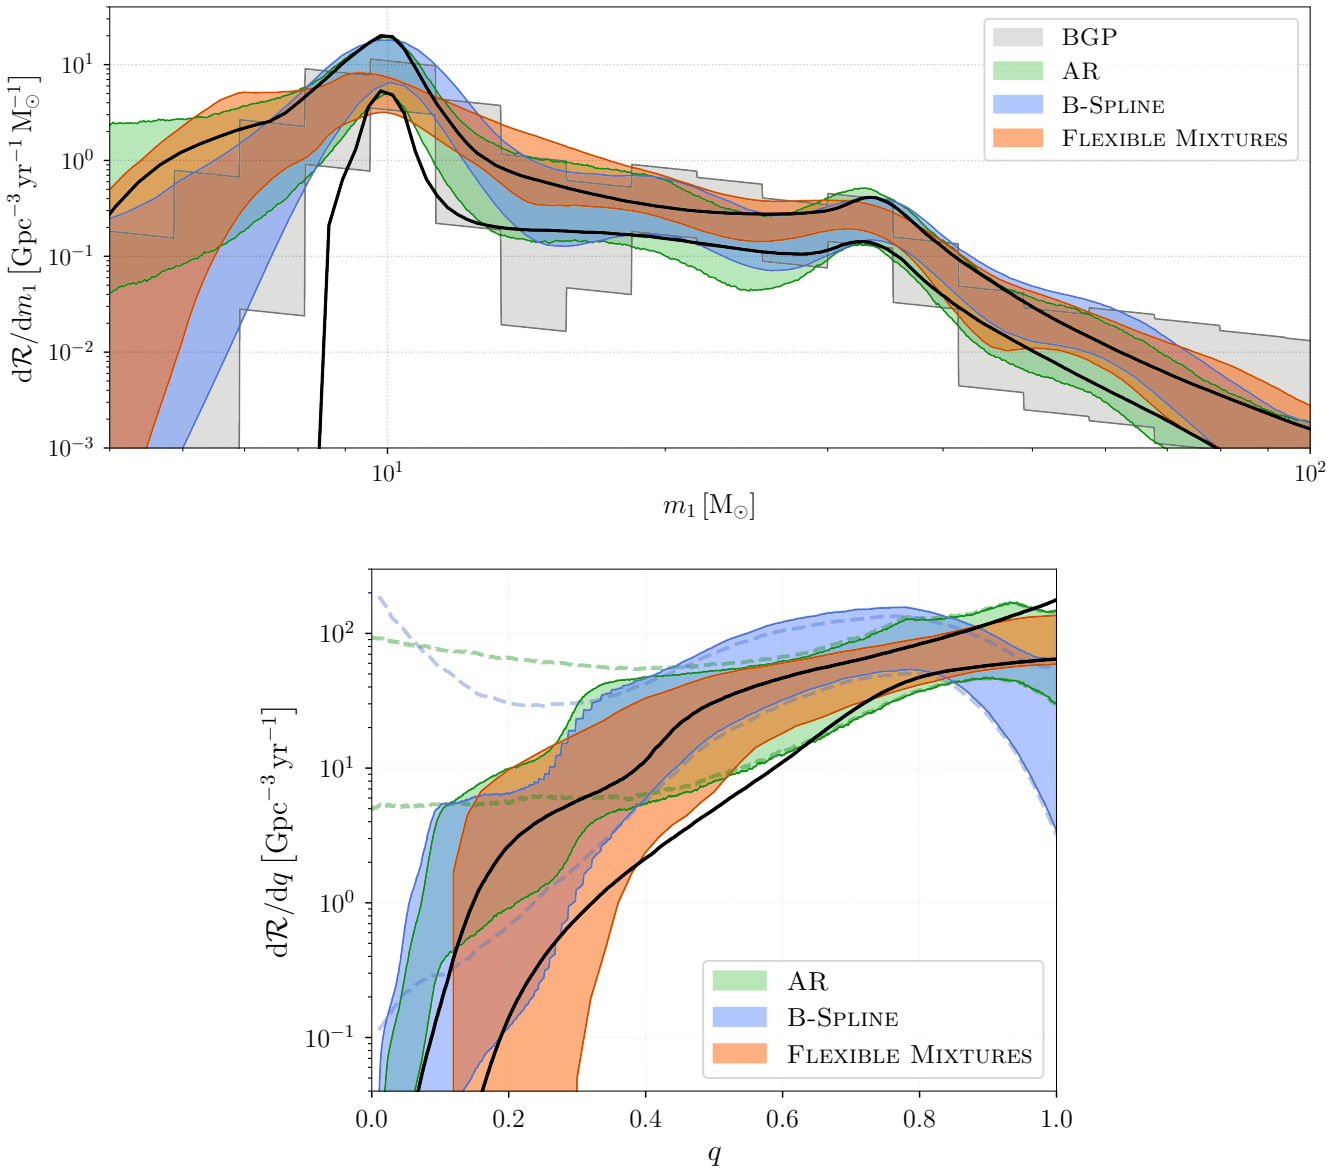

Figure 20. (*Top*) Primary mass distributions using the weakly modeled approaches outlined in Appendix C. The Broken Power Law + 2 Peaks result is shown in black for comparison. All distributions show rates evaluated at $z = 0.2$, except for BGP, which shows the rate evaluated on the $z = 0.1-0.25$ bin. Note that Flexible Mixtures does not infer primary mass directly, but instead derives it from the chirp mass Gaussian mixture model and the mass ratio as a power law distribution. The lack of substructure in the Flexible Mixtures mass distribution relative to the other models is likely due to model misspecification from assuming a power law mass ratio model. (*Bottom*) Mass ratio distributions using the weakly modeled approaches outlined in Appendix C. The Broken Power Law + 2 Peaks result is shown in black for comparison. Flexible Mixtures models the chirp mass as a Gaussian mixture model and the mass ratio as a power law. For the B-Spline and AR models, the separable distribution $p(q)$ is shown by dashed lines and the conditional marginal distribution $p(q|m_2 > 3M_\odot) = \int p(q)p(m_1)\Theta(m_1 q - 3)\mathrm{d}m_1$ is shown by the shaded regions, highlighting that the low mass ratio truncation seen in the strongly modeled approach is largely a prior effect.

- *HM fraction*: The HM fraction is a heuristic for the upper limit of the fraction of BBHs coming from hierarchical mergers; it is calculated as 0.16 times the fraction of $\chi_{\mathrm{eff}} < -0.3$ (Fishbach et al. 2022; Baibhav et al. 2020). We consistently constrain the HM fraction to be small, with the three parametric models finding it $\lesssim \mathcal{O}(10^{-1})$. The two non-parametric models, on the other hand, find it to be $\gtrsim \mathcal{O}(10^{-2})$.

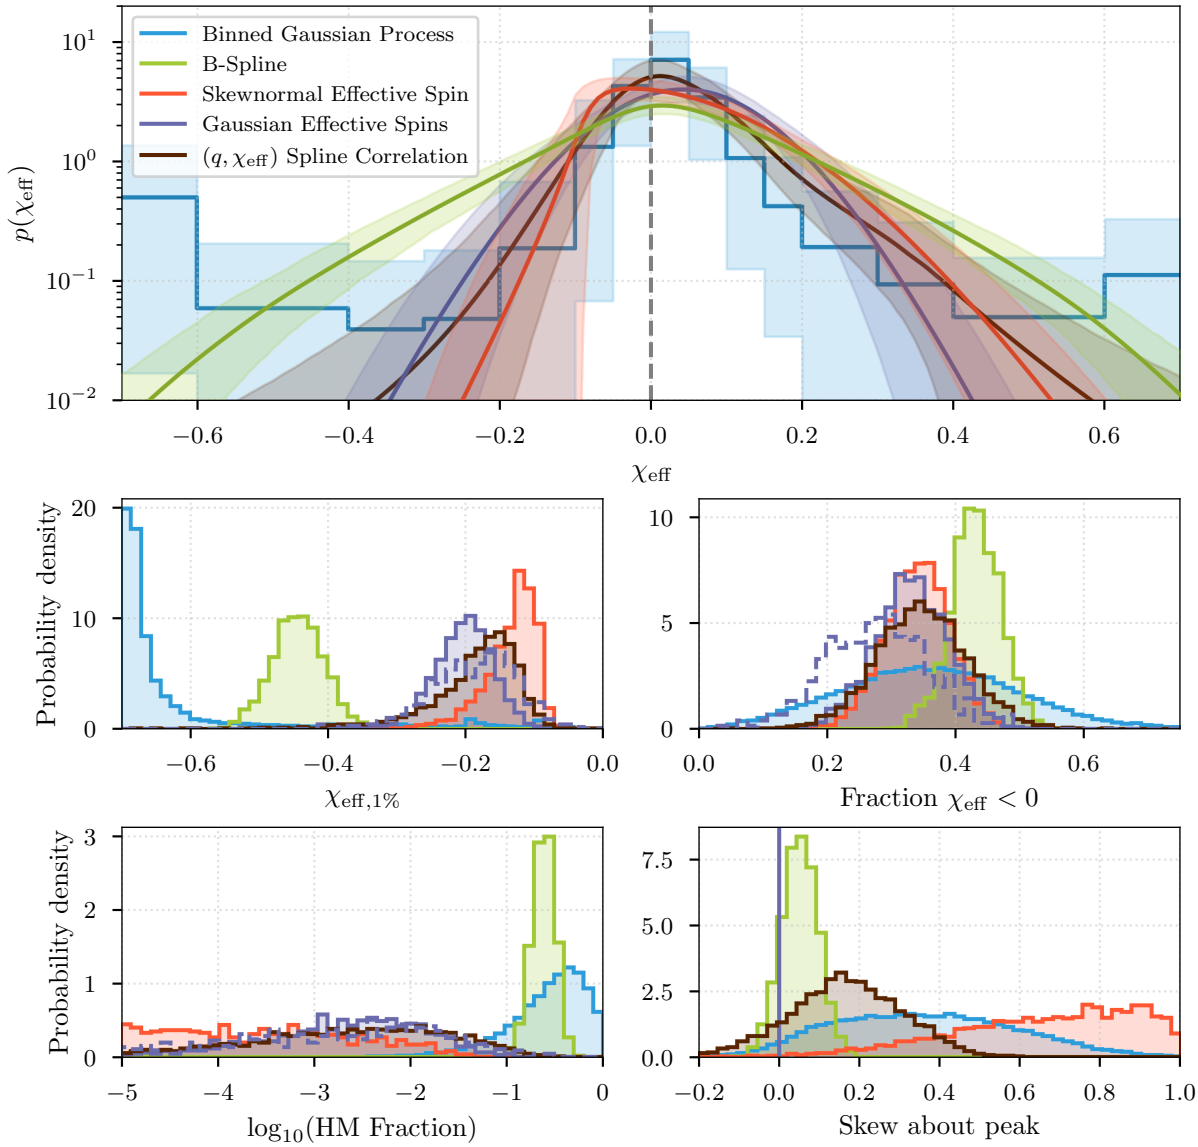

Figure 21. *First row:* Marginal $\chi_{\rm eff}$ distributions with BGP (blue) and B-Spline (green) models, as compared to the Skew-normal Effective Spin (red-orange), Gaussian Effective Spins (purple), and $(q, \chi_{\rm eff})$ Spline (maroon) models. The PPD (average) is shown with a dark line, and the 90% confidence intervals are shaded. *Second and third rows:* Posteriors on the first percentile of the $\chi_{\rm eff}$ distribution (upper left), the fraction of $\chi_{\rm eff} < 0$ (upper right), HM fraction (lower left), and the skew (lower right). See Section 6.3.2 for a discussion of these quantities. Colors correspond to the top panel; purple dashed is GWTC-3.0 with the Gaussian Effective Spins model for comparison.

- *Skew of the $\chi_{\rm eff}$ distribution about its peak*: Each model finds fairly different values of the skew, which is here defined as the difference between the fraction of $\chi_{\rm eff} > \chi_0$ and $\chi_{\rm eff} < \chi_0$, where χ_0 is the value of $\chi_{\rm eff}$ at which the distribution peaks (Banagiri et al. 2025, Equation 5). By design, the Gaussian Effective Spins model will always have a skew of ~ 0: truncated normal distributions are approximately symmetric about their peak if the truncation occurs substantially far from the distribution's bulk, which is here the case. All other models find a preference for positive skew, meaning that the distribution has more support for $\chi_{\rm eff}$ above its peak than below. This is the most pronounced for the Skew-normal Effective Spin, which is entirely inconsistent with a skew of 0.

D.5. *Merger Rates Including Subthreshold Triggers*

Search pipelines compute signal-to-noise ratios (SNRs) of detector data to identify portions of data with above-threshold SNR, i.e., "triggers", which may contain GW events (Abac et al. 2025b). Based on pipeline-specific ranking statistics, triggers are assigned significances quantified by FARs. In the main text, merger rates were calculated from population analyses that only used triggers surpassing a fixed-significance FAR threshold, which was motivated to introduce minimal contamination from noise events. By design, the list of triggers included in those analyses is threshold-dependent. Here, we explore the merger rates by including subthreshold triggers (Farr et al. 2015; Kapadia et al. 2020), which ensures that the inferred rates are free of biases due to arbitrary significance thresholds and mitigates loss of information from excluding subthreshold GW candidates. Specifically, we consider the full set of available triggers from a matched-filtering search, GstLAL (Messick et al. 2017; Sachdev et al. 2019; Hanna et al. 2020; Cannon et al. 2020; Ewing et al. 2024; Tsukada et al. 2023; Sakon et al. 2024; Ray et al. 2023b; Joshi et al. 2025a,b).

The method used by GstLAL to self-consistently classify triggers and compute $p_{\rm astro}$ (Abbott et al. 2019b, 2024, 2023b; Kapadia et al. 2020; Abac et al. 2025b,c; Ray et al. 2023b) values provides the rate densities described here for BNSs, NSBHs, and BBHs (Abac et al. 2025b; Kapadia et al. 2020). The approach is a simplified version of the methods described in Appendix C 3 of Abbott et al. (2023a), as well as the discovery of GW200105 and GW200115 (Abbott et al. 2021c) and GW230529 (Abac et al. 2024). The simplification involves fixing the mass distribution to the Salpeter model (Salpeter 1955) as opposed to marginalizing over population uncertainties inferred by the above-threshold event analyses explored in the main text. We do not expect the marginalization over population uncertainties to have a significant impact on the inferred merger rates, given the uncertainty ranges. When we apply the simpler model to GWTC-3.0 (Abbott et al. 2023a), the estimated rates are comparable to those derived when marginalizing over the inferred population. For effective inspiral spins of each trigger, we use a uniform distribution. Consistent with how searches categorize triggers and compute $\langle VT\rangle$ (Abac et al. 2025b), we set the boundary between NSs and BHs to be 3 $M_\odot$. For triggers categorized as BNS or NSBH, we set the bounds of the spins corresponding to a NS to ± 0.4. These bounds are consistent with what the analyses in the main text adopt (see Section 3.2 and Section 5).

Using the fixed population distribution and the Poisson mixture model of Kapadia et al. (2020), we infer the merger rate densities of the different source categories (BNS, NSBH, and BBH) from the full list of available GstLAL triggers with FAR < 1 hour^{-1}. This is done while self-consistently accounting for the possibility that some of these triggers are likely noise artifacts. We construct the posterior of astrophysical counts of BNS, NSBH, and BBH events by utilizing mass-based binning template weights (Ray et al. 2023b). We estimate the time–volume sensitivity (Tiwari 2018; Abac et al. 2025b) $\langle VT\rangle$ for each event category α (α = BNS, NSBH, BBH), using injections (Essick et al. 2025), and their contributions to the counts posterior (Kapadia et al. 2020). Finally, we compute the rates posterior from marginalized counts posterior and the estimated $\langle VT\rangle$ (Abbott et al. 2019b, 2021a, 2023a; Abac et al. 2025c):

$$p\left(R_\alpha\right) = p\left(\Lambda_{1\alpha}|\mathbf{x}\right)\langle VT\rangle_{\text{O1-O4a},\alpha}. \tag{D46}$$

Here, $p\left(\Lambda_{1\alpha}|\mathbf{x}\right)$ is the marginalized counts posterior where $\Lambda_{1\alpha}$ is the astrophysical count for event category α and $\Lambda_{1\alpha} = R_\alpha\langle VT\rangle_{\text{O1-O4a},\alpha}$. Here, O1-O4a indicates that data from O1 throughout O4a are used. The vector $\mathbf{x}$ represents the set of triggers using in this analysis, where each x_i consists of the ranking statistics information, SNR, and an identifier for the template associated with the trigger.

A Jeffreys prior, $\propto N^{-1/2}$, is imposed on the astrophysical counts for BBHs to construct their posterior from the mixed Poisson likelihood and compute the merger rate. The merger rate of BBHs is computed to be 13.1–17.3 Gpc^{-3} yr^{-1}. This is consistent with the BBH merger rate provided in the analyses presented in the main text, and the rate has been further narrowed compared to the results of Abbott et al. (2023a) while still being consistent within uncertainties. A uniform prior is used for NSBHs and BNSs, as the number of detected events containing a NS is small. We compute the NSBH merger rates to be 22.2–143 Gpc^{-3} yr^{-1}, which is consistent with the NSBH merger rate presented in main text. Similarly to the rate for BBHs, the rate of NSBHs is consistent with the rate obtained from the analyses of GWTC-3.0 but has been further narrowed. The BNS merger rate is calculated to be 30.9–361 Gpc^{-3} yr^{-1}, which is consistent with the BNS merger rate presented in the main text and has been narrowed down compared to the rate obtained from the analyses of GWTC-3.0.

D.6. *Supplementary Results: BBH Correlations*

Here, we supplement the correlated BBH population results presented in Section 6.5 with additional results and figures. In the top panel of Figure 22, we plot the posterior distribution for the level of correlation $\kappa_{\chi_{\rm eff},z}$ inferred in the COPULA model for $(z, \chi_{\rm eff})$. Note the peak at positive values, and long tail into negative values. The middle panels show the inferred mass distributions, using a mass-redshift correlated BGP analysis, binned by redshift. We see that both the primary and secondary mass distribution appear to be broadly consistent across redshifts up to $z = 1$. Finally, we investigate mass and spin correlations with the data-driven FM analyses. In the bottom panels of Figure 22, we plot the inferred distribution of orbital aligned spin χ_z as a function of both chirp mass and mass ratio, obtained with the FM analysis. We see that the results have a large degree of uncertainty, and are broadly consistent with findings explored here and in Section 6.5.

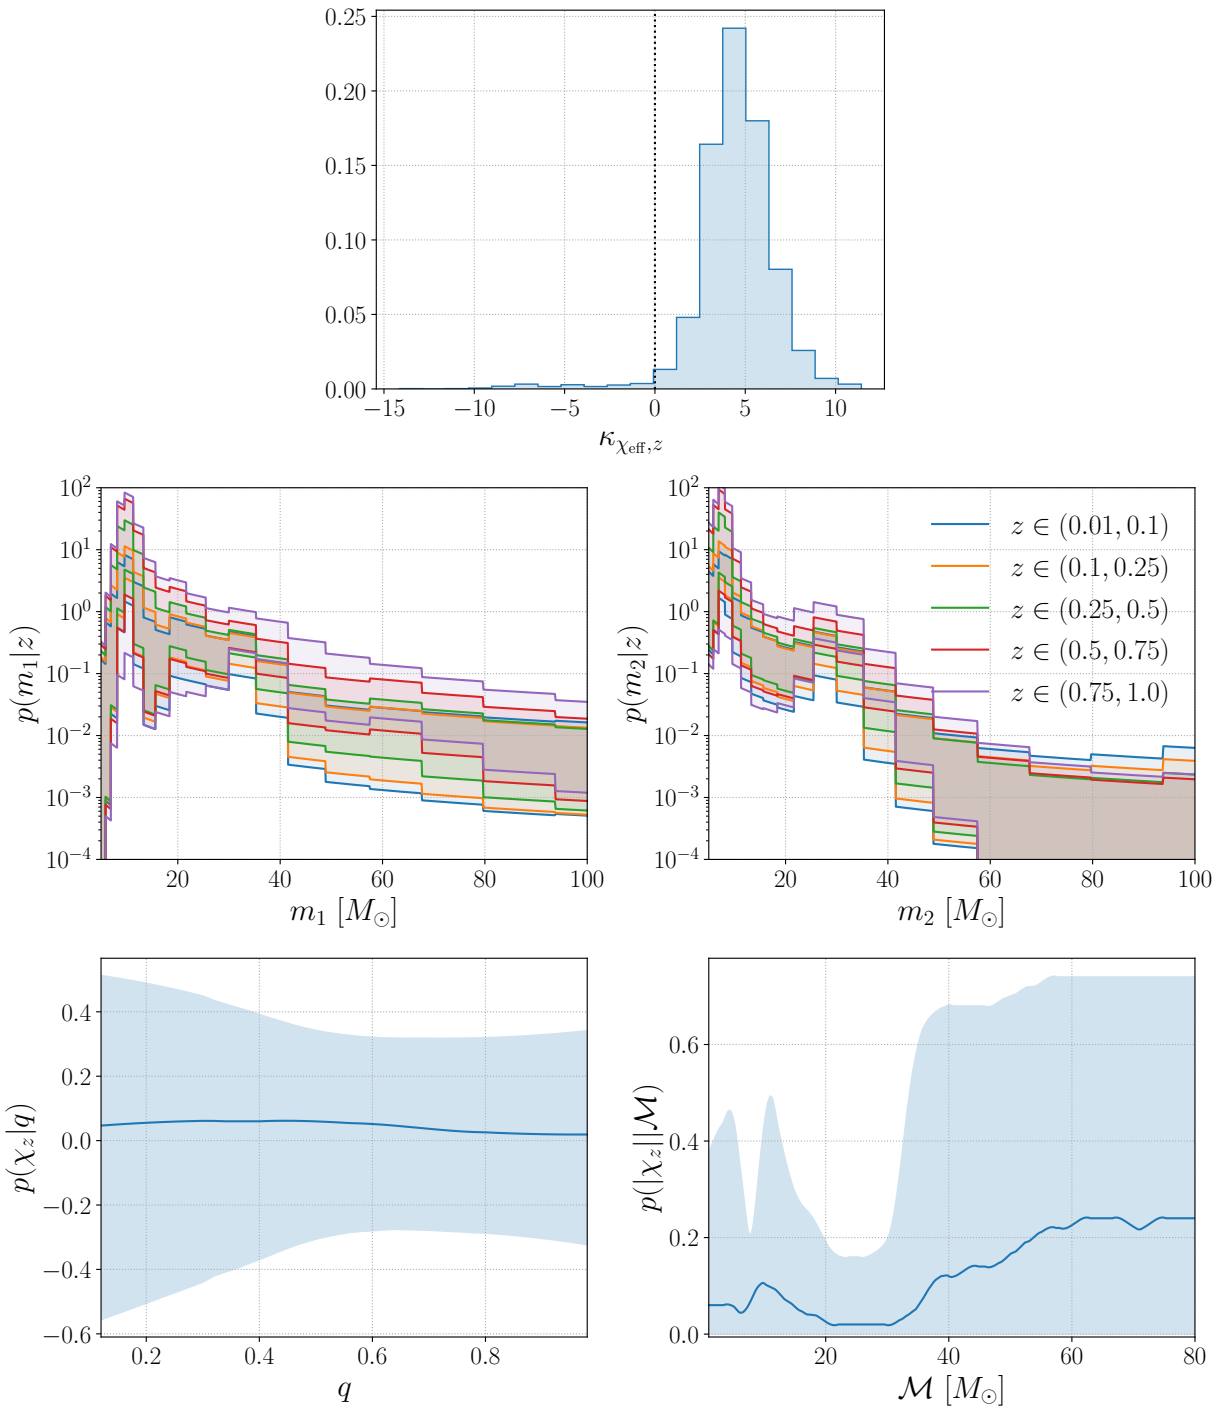

Figure 22. *Top:* Posterior distributions for the level of correlation between redshift and effective inspiral spin $\kappa_{\chi_{\rm eff},z}$ inferred using the COPULA model. The vertical black dashed line in each plot indicates a value of $\kappa_{\chi_{\rm eff},z}$, at which no correlation is implied. *Middle:* Inferred distributions of primary mass (left), and secondary mass (right) in the redshift and spin correlated BGP analysis. The solid lines bound the 90% credible intervals for each redshift bin. We can see that all redshift bins from $z = 0.01$ to $z = 1$ are consistent within 90% credibility. *Bottom:* Correlated mass and spin PPDs from the FM model. Solid lines give the medians, while the shaded regions encompass 90% of the PPD volume. The left panel gives the inferred distribution of aligned spin χ_z given mass ratio. We do not see any evidence for or against a correlation. The right panel gives the inferred distribution of the aligned spin magnitude $|\chi_z|$ as a function of chirp mass. We see that the uncertainty is very large, with a notable drop in the $\sim$20$-$30 $M_\odot$ region. For reference, this region very roughly corresponds to the 30 $M_\odot$ peak observed in the component mass distributions (a 30 $M_\odot$ + 30 $M_\odot$ BBH has a chirp mass of $\mathcal{M} \approx 26\ M_\odot$). Following the mass- and spin-correlated BGP results then, it is unsurprising that the FM results prefer a local minimum in aligned spin magnitude in this region.

REFERENCES

Aasi, J., et al. 2015, Class. Quant. Grav., 32, 074001, doi: 10.1088/0264-9381/32/7/074001

Abac, A. G., et al. 2024, Astrophys. J. Lett., 970, L34, doi: 10.3847/2041-8213/ad5beb

—. 2025a, To be published in this issue. https://dcc.ligo.org/LIGO-P2400293/public

—. 2025b, To be published in this issue. https://dcc.ligo.org/LIGO-P2400300/public

—. 2025c, To be published in this issue. https://dcc.ligo.org/LIGO-P2400386/public

—. 2025d, To be published in this issue. https://dcc.ligo.org/LIGO-P2500167/public

—. 2025e. https://arxiv.org/abs/2507.08219

Abbott, B. P., et al. 2016, Phys. Rev. X, 6, 041015, doi: 10.1103/PhysRevX.6.041015

—. 2017, Phys. Rev. Lett., 119, 161101, doi: 10.1103/PhysRevLett.119.161101

—. 2019a, Astrophys. J. Lett., 882, L24, doi: 10.3847/2041-8213/ab3800

—. 2019b, Phys. Rev. X, 9, 031040, doi: 10.1103/PhysRevX.9.031040

—. 2019c, Phys. Rev. X, 9, 011001, doi: 10.1103/PhysRevX.9.011001

—. 2020a, Living Rev. Rel., 23, 3, doi: 10.1007/s41114-020-00026-9

—. 2020b, Class. Quant. Grav., 37, 045006, doi: 10.1088/1361-6382/ab5f7c

Abbott, R., et al. 2020c, Phys. Rev. Lett., 125, 101102, doi: 10.1103/PhysRevLett.125.101102

—. 2020d, Astrophys. J. Lett., 900, L13, doi: 10.3847/2041-8213/aba493

—. 2020e, Astrophys. J. Lett., 896, L44, doi: 10.3847/2041-8213/ab960f

—. 2021a, Astrophys. J. Lett., 913, L7, doi: 10.3847/2041-8213/abe949

—. 2021b, Phys. Rev. X, 11, 021053, doi: 10.1103/PhysRevX.11.021053

—. 2021c, Astrophys. J. Lett., 915, L5, doi: 10.3847/2041-8213/ac082e

—. 2023a, Phys. Rev. X, 13, 011048, doi: 10.1103/PhysRevX.13.011048

—. 2023b, Phys. Rev. X, 13, 041039, doi: 10.1103/PhysRevX.13.041039

—. 2024, Phys. Rev. D, 109, 022001, doi: 10.1103/PhysRevD.109.022001

Acernese, F., et al. 2015, Class. Quant. Grav., 32, 024001, doi: 10.1088/0264-9381/32/2/024001

Adamcewicz, C., Lasky, P. D., & Thrane, E. 2023, Astrophys. J., 958, 13, doi: 10.3847/1538-4357/acf763

Adamcewicz, C., & Thrane, E. 2022, Mon. Not. Roy. Astron. Soc., 517, 3928, doi: 10.1093/mnras/stac2961

Adams, T., Buskulic, D., Germain, V., et al. 2016, Class. Quant. Grav., 33, 175012, doi: 10.1088/0264-9381/33/17/175012

Afroz, S., & Mukherjee, S. 2025, Phys. Rev. D, 112, 023531, doi: 10.1103/7zc2-g9vq

Ai, S., Gao, H., Yuan, Y., Zhang, B., & Lan, L. 2023, Mon. Not. Roy. Astron. Soc., 526, 6260, doi: 10.1093/mnras/stad3177

Ajith, P., et al. 2011, Phys. Rev. Lett., 106, 241101, doi: 10.1103/PhysRevLett.106.241101

Akutsu, T., et al. 2021, PTEP, 2021, 05A101, doi: 10.1093/ptep/ptaa125

Alléné, C., et al. 2025, Class. Quant. Grav., 42, 105009, doi: 10.1088/1361-6382/add234

Andres, N., et al. 2022, Class. Quant. Grav., 39, 055002, doi: 10.1088/1361-6382/ac482a

Antonini, F., Callister, T., Dosopoulou, F., Romero-Shaw, I., & Chattopadhyay, D. 2025a. https://arxiv.org/abs/2506.09154

Antonini, F., & Gieles, M. 2020, Phys. Rev. D, 102, 123016, doi: 10.1103/PhysRevD.102.123016

Antonini, F., Gieles, M., & Gualandris, A. 2019, Mon. Not. Roy. Astron. Soc., 486, 5008, doi: 10.1093/mnras/stz1149

Antonini, F., Romero-Shaw, I. M., & Callister, T. 2025b, Phys. Rev. Lett., 134, 011401, doi: 10.1103/PhysRevLett.134.011401

Arca Sedda, M. 2020, Astrophys. J., 891, 47, doi: 10.3847/1538-4357/ab723b

Arca Sedda, M., Kamlah, A. W. H., Spurzem, R., et al. 2024, Mon. Not. Roy. Astron. Soc., 528, 5140, doi: 10.1093/mnras/stad3951

Arca Sedda, M., Kamlah, A. W. H., Spurzem, R., et al. 2023, MNRAS, 526, 429, doi: 10.1093/mnras/stad2292

Arun, K. G., Buonanno, A., Faye, G., & Ochsner, E. 2009, Phys. Rev. D, 79, 104023, doi: 10.1103/PhysRevD.79.104023

Ashton, G., et al. 2019, Astrophys. J. Suppl., 241, 27, doi: 10.3847/1538-4365/ab06fc

Aubin, F., et al. 2021, Class. Quant. Grav., 38, 095004, doi: 10.1088/1361-6382/abe913

Baibhav, V., Berti, E., Gerosa, D., et al. 2019, Phys. Rev. D, 100, 064060, doi: 10.1103/PhysRevD.100.064060

Baibhav, V., Doctor, Z., & Kalogera, V. 2023, Astrophys. J., 946, 50, doi: 10.3847/1538-4357/acbf4c

Baibhav, V., Gerosa, D., Berti, E., et al. 2020, Phys. Rev. D, 102, 043002, doi: 10.1103/PhysRevD.102.043002

Baibhav, V., & Kalogera, V. 2024. https://arxiv.org/abs/2412.03461
Bailyn, C. D., Jain, R. K., Coppi, P., & Orosz, J. A. 1998, Astrophys. J., 499, 367, doi: 10.1086/305614
Banagiri, S., Callister, T. A., Doctor, Z., & Kalogera, V. 2025. https://arxiv.org/abs/2501.06712
Banerjee, S. 2021, Mon. Not. Roy. Astron. Soc., 503, 3371, doi: 10.1093/mnras/stab591
Banerjee, S., Olejak, A., & Belczynski, K. 2023, Astrophys. J., 953, 80, doi: 10.3847/1538-4357/acdd59
Barr, E. D., et al. 2024, Science, 383, 275, doi: 10.1126/science.adg3005
Bartos, I., Kocsis, B., Haiman, Z., & Márka, S. 2017, Astrophys. J., 835, 165, doi: 10.3847/1538-4357/835/2/165
Bavera, S. S., Fishbach, M., Zevin, M., Zapartas, E., & Fragos, T. 2022, Astron. Astrophys., 665, A59, doi: 10.1051/0004-6361/202243724
Bavera, S. S., Fragos, T., Qin, Y., et al. 2020, Astron. Astrophys., 635, A97, doi: 10.1051/0004-6361/201936204
Bavera, S. S., et al. 2021, Astron. Astrophys., 647, A153, doi: 10.1051/0004-6361/202039804
Belczynski, K., Kalogera, V., & Bulik, T. 2001, Astrophys. J., 572, 407, doi: 10.1086/340304
Belczynski, K., Wiktorowicz, G., Fryer, C., Holz, D., & Kalogera, V. 2012, Astrophys. J., 757, 91, doi: 10.1088/0004-637X/757/1/91
Belczynski, K., et al. 2020, Astron. Astrophys., 636, A104, doi: 10.1051/0004-6361/201936528
Bethe, H. A., & Brown, G. E. 1998, Astrophys. J., 506, 780, doi: 10.1086/306265
Bianconi, M., Smith, G. P., Nicholl, M., et al. 2023, Mon. Not. Roy. Astron. Soc., 521, 3421, doi: 10.1093/mnras/stad673
Bingham, E., Chen, J. P., Jankowiak, M., et al. 2019, J. Mach. Learn. Res., 20, 28:1. http://jmlr.org/papers/v20/18-403.html
Bird, S., Cholis, I., Muñoz, J. B., et al. 2016, Phys. Rev. Lett., 116, 201301, doi: 10.1103/PhysRevLett.116.201301
Biscoveanu, S., Callister, T. A., Haster, C.-J., et al. 2022a, Astrophys. J. Lett., 932, L19, doi: 10.3847/2041-8213/ac71a8
Biscoveanu, S., Isi, M., Vitale, S., & Varma, V. 2021, Phys. Rev. Lett., 126, 171103, doi: 10.1103/PhysRevLett.126.171103
Biscoveanu, S., Landry, P., & Vitale, S. 2022b, Mon. Not. Roy. Astron. Soc., 518, 5298, doi: 10.1093/mnras/stac3052
Boesky, A. P., Broekgaarden, F. S., & Berger, E. 2024, Astrophys. J., 976, 24, doi: 10.3847/1538-4357/ad7fe3
Bradbury, J., Frostig, R., Hawkins, P., et al. 2018, JAX: composable transformations of Python+NumPy programs, 0.3.13. http://github.com/google/jax
Broekgaarden, F. S., Stevenson, S., & Thrane, E. 2022a, Astrophys. J., 938, 45, doi: 10.3847/1538-4357/ac8879
Broekgaarden, F. S., et al. 2022b, Mon. Not. Roy. Astron. Soc., 516, 5737, doi: 10.1093/mnras/stac1677
Burrows, A., & Vartanyan, D. 2021, Nature, 589, 29, doi: 10.1038/s41586-020-03059-w
Callister, T. A., Essick, R., & Holz, D. E. 2024, Phys. Rev. D, 110, 123041, doi: 10.1103/PhysRevD.110.123041
Callister, T. A., & Farr, W. M. 2024, Phys. Rev. X, 14, 021005, doi: 10.1103/PhysRevX.14.021005
Callister, T. A., Farr, W. M., & Renzo, M. 2021a, Astrophys. J., 920, 157, doi: 10.3847/1538-4357/ac1347
Callister, T. A., Haster, C.-J., Ng, K. K. Y., Vitale, S., & Farr, W. M. 2021b, Astrophys. J. Lett., 922, L5, doi: 10.3847/2041-8213/ac2ccc
Callister, T. A., Miller, S. J., Chatziioannou, K., & Farr, W. M. 2022, Astrophys. J. Lett., 937, L13, doi: 10.3847/2041-8213/ac847e
Cannon, K., et al. 2020. https://arxiv.org/abs/2010.05082
Capote, E., et al. 2025, Phys. Rev. D, 111, 062002, doi: 10.1103/PhysRevD.111.062002
Cheng, A. Q., Zevin, M., & Vitale, S. 2023, Astrophys. J., 955, 127, doi: 10.3847/1538-4357/aced98
Chruślińska, M. 2024, Annalen Phys., 536, 2200170, doi: 10.1002/andp.202200170
Clesse, S., & Garcia-Bellido, J. 2022, Phys. Dark Univ., 38, 101111, doi: 10.1016/j.dark.2022.101111
Colleoni, M., Vidal, F. A. R., García-Quirós, C., Akçay, S., & Bera, S. 2025, Phys. Rev. D, 111, 104019, doi: 10.1103/PhysRevD.111.104019
Colloms, S., Berry, C. P. L., Veitch, J., & Zevin, M. 2025, Astrophys. J., 988, 189, doi: 10.3847/1538-4357/ade546
Dal Canton, T., Nitz, A. H., Gadre, B., et al. 2021, Astrophys. J., 923, 254, doi: 10.3847/1538-4357/ac2f9a
Damour, T. 2001, Phys. Rev. D, 64, 124013, doi: 10.1103/PhysRevD.64.124013
de Mink, S. E., & Mandel, I. 2016, Mon. Not. Roy. Astron. Soc., 460, 3545, doi: 10.1093/mnras/stw1219
de Sá, L. M., Rocha, L. S., Bernardo, A., Bachega, R. R. A., & Horvath, J. E. 2024, Mon. Not. Roy. Astron. Soc., 535, 2041, doi: 10.1093/mnras/stae2281
Delfavero, V., O'Shaughnessy, R., Wysocki, D., & Yelikar, A. 2021. https://arxiv.org/abs/2107.13082
Di Carlo, U. N., Giacobbo, N., Mapelli, M., et al. 2019, Mon. Not. Roy. Astron. Soc., 487, 2947, doi: 10.1093/mnras/stz1453

Di Carlo, U. N., Mapelli, M., Bouffanais, Y., et al. 2020, Mon. Not. Roy. Astron. Soc., 497, 1043, doi: 10.1093/mnras/staa1997
Dietrich, T., Coughlin, M. W., Pang, P. T. H., et al. 2020, Science, 370, 1450, doi: 10.1126/science.abb4317
Dietrich, T., Samajdar, A., Khan, S., et al. 2019, Phys. Rev. D, 100, 044003, doi: 10.1103/PhysRevD.100.044003
Dittmann, A. J., et al. 2024, Astrophys. J., 974, 295, doi: 10.3847/1538-4357/ad5f1e
Doctor, Z., Wysocki, D., O'Shaughnessy, R., Holz, D. E., & Farr, B. 2019, doi: 10.3847/1538-4357/ab7fac
Dominik, M., Belczynski, K., Fryer, C., et al. 2013, Astrophys. J., 779, 72, doi: 10.1088/0004-637X/779/1/72
Dominik, M., Berti, E., O'Shaughnessy, R., et al. 2015, Astrophys. J., 806, 263, doi: 10.1088/0004-637X/806/2/263
Drago, M., et al. 2020, doi: 10.1016/j.softx.2021.100678
du Buisson, L., Marchant, P., Podsiadlowski, P., et al. 2020, Mon. Not. Roy. Astron. Soc., 499, 5941, doi: 10.1093/mnras/staa3225
Edelman, B., Doctor, Z., Godfrey, J., & Farr, B. 2022, Astrophys. J., 924, 101, doi: 10.3847/1538-4357/ac3667
Edelman, B., Farr, B., & Doctor, Z. 2023, Astrophys. J., 946, 16, doi: 10.3847/1538-4357/acb5ed
El-Badry, K., Rix, H.-W., Latham, D. W., et al. 2024, The Open Journal of Astrophysics, 7, 58, doi: 10.33232/001c.121261
Ertl, T., Woosley, S. E., Sukhbold, T., & Janka, H. T. 2019, doi: 10.3847/1538-4357/ab6458
Essick, R. 2025a, GWTC-4: O4a Search Sensitivity Estimates, Zenodo, doi: 10.5281/zenodo.16740117
—. 2025b, GWTC-4: Cumulative Search Sensitivity Estimates, Zenodo, doi: 10.5281/zenodo.16740128
Essick, R., Farah, A., Galaudage, S., et al. 2022, Astrophys. J., 926, 34, doi: 10.3847/1538-4357/ac3978
Essick, R., & Farr, W. 2022. https://arxiv.org/abs/2204.00461
Essick, R., & Fishbach, M. 2024, Astrophys. J., 962, 169, doi: 10.3847/1538-4357/ad1604
Essick, R., et al. 2025. https://arxiv.org/abs/2508.10638
Ewing, B., et al. 2024, Phys. Rev. D, 109, 042008, doi: 10.1103/PhysRevD.109.042008
Fairhurst, S., Green, R., Hoy, C., Hannam, M., & Muir, A. 2020, Phys. Rev. D, 102, 024055, doi: 10.1103/PhysRevD.102.024055
Farah, A., Ezquiaga, J. M., Fishbach, M., & Holz, D. 2025a. https://arxiv.org/abs/2507.07964
Farah, A. M., Callister, T. A., Ezquiaga, J. M., Zevin, M., & Holz, D. E. 2025b, Astrophys. J., 978, 153, doi: 10.3847/1538-4357/ad9253
Farah, A. M., Fishbach, M., Essick, R., Holz, D. E., & Galaudage, S. 2022, Astrophys. J., 931, 108, doi: 10.3847/1538-4357/ac5f03
Farmer, R., Renzo, M., de Mink, S., Fishbach, M., & Justham, S. 2020, Astrophys. J. Lett., 902, L36, doi: 10.3847/2041-8213/abbadd
Farmer, R., Renzo, M., de Mink, S. E., Marchant, P., & Justham, S. 2019, doi: 10.3847/1538-4357/ab518b
Farr, B., & Farr, W. 2025, On the incompatability of q–$\chi_{\rm eff}$ correlations and IID spins, LIGO DCC. https://dcc.ligo.org/T2500277
Farr, W. M. 2019, Research Notes of the AAS, 3, 66, doi: 10.3847/2515-5172/ab1d5f
Farr, W. M., Gair, J. R., Mandel, I., & Cutler, C. 2015, Phys. Rev. D, 91, 023005, doi: 10.1103/PhysRevD.91.023005
Farr, W. M., Sravan, N., Cantrell, A., et al. 2011, Astrophys. J., 741, 103, doi: 10.1088/0004-637X/741/2/103
Farr, W. M., Stevenson, S., Coleman Miller, M., et al. 2017, Nature, 548, 426, doi: 10.1038/nature23453
Farrow, N., Zhu, X.-J., & Thrane, E. 2019, Astrophys. J., 876, 18, doi: 10.3847/1538-4357/ab12e3
Fishbach, M., Essick, R., & Holz, D. E. 2020, Astrophys. J. Lett., 899, L8, doi: 10.3847/2041-8213/aba7b6
Fishbach, M., & Fragione, G. 2023, Mon. Not. Roy. Astron. Soc., 522, 5546, doi: 10.1093/mnras/stad1364
Fishbach, M., Holz, D. E., & Farr, B. 2017, Astrophys. J. Lett., 840, L24, doi: 10.3847/2041-8213/aa7045
Fishbach, M., Holz, D. E., & Farr, W. M. 2018, Astrophys. J. Lett., 863, L41, doi: 10.3847/2041-8213/aad800
Fishbach, M., & Kalogera, V. 2021, Astrophys. J. Lett., 914, L30, doi: 10.3847/2041-8213/ac05c4
Fishbach, M., Kimball, C., & Kalogera, V. 2022, Astrophys. J. Lett., 935, L26, doi: 10.3847/2041-8213/ac86c4
Fishbach, M., & van Son, L. 2023, Astrophys. J. Lett., 957, L31, doi: 10.3847/2041-8213/ad0560
Fishbach, M., Doctor, Z., Callister, T., et al. 2021, Astrophys. J., 912, 98, doi: 10.3847/1538-4357/abee11
Fragione, G., Grishin, E., Leigh, N. W. C., Perets, H. B., & Perna, R. 2019, Mon. Not. Roy. Astron. Soc., 488, 47, doi: 10.1093/mnras/stz1651
Fragione, G., & Kocsis, B. 2018, Phys. Rev. Lett., 121, 161103, doi: 10.1103/PhysRevLett.121.161103
Fragione, G., & Silk, J. 2020, Mon. Not. Roy. Astron. Soc., 498, 4591, doi: 10.1093/mnras/staa2629
Fryer, C. L., Belczynski, K., Wiktorowicz, G., et al. 2012, Astrophys. J., 749, 91, doi: 10.1088/0004-637X/749/1/91
Fryer, C. L., & Kalogera, V. 2001, Astrophys. J., 554, 548, doi: 10.1086/321359

Fryer, C. L., Olejak, A., & Belczynski, K. 2022, Astrophys. J., 931, 94, doi: 10.3847/1538-4357/ac6ac9

Fuller, J., & Lu, W. 2022, Mon. Not. Roy. Astron. Soc., 511, 3951, doi: 10.1093/mnras/stac317

Fuller, J., & Ma, L. 2019, Astrophys. J. Lett., 881, L1, doi: 10.3847/2041-8213/ab339b

Fuller, J., Piro, A. L., & Jermyn, A. S. 2019, Mon. Not. Roy. Astron. Soc., 485, 3661, doi: 10.1093/mnras/stz514

Galaudage, S., & Lamberts, A. 2025, Astron. Astrophys., 694, A186, doi: 10.1051/0004-6361/202451654

Galaudage, S., et al. 2021, Astrophys. J. Lett., 921, L15, doi: 10.3847/2041-8213/ac2f3c

Gallegos-Garcia, M., Berry, C. P. L., Marchant, P., & Kalogera, V. 2021, Astrophys. J., 922, 110, doi: 10.3847/1538-4357/ac2610

Ganapathy, D., et al. 2023, Phys. Rev. X, 13, 041021, doi: 10.1103/PhysRevX.13.041021

Gerosa, D., & Berti, E. 2017, Phys. Rev. D, 95, 124046, doi: 10.1103/PhysRevD.95.124046

Gerosa, D., Berti, E., O'Shaughnessy, R., et al. 2018, Phys. Rev. D, 98, 084036, doi: 10.1103/PhysRevD.98.084036

Gerosa, D., & Fishbach, M. 2021, Nature Astron., 5, 749, doi: 10.1038/s41550-021-01398-w

Gerosa, D., Kesden, M., Sperhake, U., Berti, E., & O'Shaughnessy, R. 2015, Phys. Rev. D, 92, 064016, doi: 10.1103/PhysRevD.92.064016

Gerosa, D., Mould, M., Gangardt, D., et al. 2021, Phys. Rev. D, 103, 064067, doi: 10.1103/PhysRevD.103.064067

Gerosa, D., Pratten, G., & Vecchio, A. 2020, Phys. Rev. D, 102, 103020, doi: 10.1103/PhysRevD.102.103020

Giacobbo, N., & Mapelli, M. 2018, Mon. Not. Roy. Astron. Soc., 480, 2011, doi: 10.1093/mnras/sty1999

Godfrey, J., Edelman, B., & Farr, B. 2023. https://arxiv.org/abs/2304.01288

Golomb, J., & Talbot, C. 2022, Astrophys. J., 926, 79, doi: 10.3847/1538-4357/ac43bc

—. 2023, Phys. Rev. D, 108, 103009, doi: 10.1103/PhysRevD.108.103009

Guo, W.-H., Li, Y.-J., Wang, Y.-Z., et al. 2024, Astrophys. J., 975, 54, doi: 10.3847/1538-4357/ad758a

Gupta, A., Gerosa, D., Arun, K. G., et al. 2020, Phys. Rev. D, 101, 103036, doi: 10.1103/PhysRevD.101.103036

Hanna, C., et al. 2020, Phys. Rev. D, 101, 022003, doi: 10.1103/PhysRevD.101.022003

Heinzel, J., Biscoveanu, S., & Vitale, S. 2024, Phys. Rev. D, 109, 103006, doi: 10.1103/PhysRevD.109.103006

Heinzel, J., Mould, M., Álvarez-López, S., & Vitale, S. 2025a, Phys. Rev. D, 111, 063043, doi: 10.1103/PhysRevD.111.063043

Heinzel, J., Mould, M., & Vitale, S. 2025b, Phys. Rev. D, 111, L061305, doi: 10.1103/PhysRevD.111.L061305

Hendriks, D. D., van Son, L. A. C., Renzo, M., Izzard, R. G., & Farmer, R. 2023, Mon. Not. Roy. Astron. Soc., 526, 4130, doi: 10.1093/mnras/stad2857

Hong, J., Vesperini, E., Askar, A., et al. 2018, Mon. Not. Roy. Astron. Soc., 480, 5645, doi: 10.1093/mnras/sty2211

Hurley, J. R., Tout, C. A., & Pols, O. R. 2002, Mon. Not. Roy. Astron. Soc., 329, 897, doi: 10.1046/j.1365-8711.2002.05038.x

Hussain, A., Isi, M., & Zimmerman, A. 2024. https://arxiv.org/abs/2411.02252

Hut, P. 1981, A&A, 99, 126

Huth, S., et al. 2022, Nature, 606, 276, doi: 10.1038/s41586-022-04750-w

Iwaya, M., Kobayashi, K., Morisaki, S., Hotokezaka, K., & Kinugawa, T. 2025, Phys. Rev. D, 111, 103046, doi: 10.1103/PhysRevD.111.103046

Janquart, J., et al. 2024, doi: 10.1093/mnras/staf049

Jayasinghe, T., et al. 2021, Mon. Not. Roy. Astron. Soc., 504, 2577, doi: 10.1093/mnras/stab907

Joshi, P., et al. 2025a. https://arxiv.org/abs/2506.06497

—. 2025b. https://arxiv.org/abs/2505.23959

Kalogera, V. 2000, Astrophys. J., 541, 319, doi: 10.1086/309400

Kalogera, V., & Baym, G. 1996, Astrophys. J. Lett., 470, L61, doi: 10.1086/310296

Kapadia, S. J., et al. 2020, Class. Quant. Grav., 37, 045007, doi: 10.1088/1361-6382/ab5f2d

Karathanasis, C., Mukherjee, S., & Mastrogiovanni, S. 2023, Mon. Not. Roy. Astron. Soc., 523, 4539, doi: 10.1093/mnras/stad1373

Kimball, C., Talbot, C., L. Berry, C. P., et al. 2020, Astrophys. J., 900, 177, doi: 10.3847/1538-4357/aba518

Kıroğlu, F., Lombardi, J. C., Kremer, K., Vanderzyden, H. D., & Rasio, F. A. 2025, Astrophys. J. Lett., 983, L9, doi: 10.3847/2041-8213/adc263

Klimenko, S. 2022. https://arxiv.org/abs/2201.01096

Klimenko, S., Mohanty, S., Rakhmanov, M., & Mitselmakher, G. 2005, Phys. Rev. D, 72, 122002, doi: 10.1103/PhysRevD.72.122002

Klimenko, S., Yakushin, I., Mercer, A., & Mitselmakher, G. 2008, Class. Quant. Grav., 25, 114029, doi: 10.1088/0264-9381/25/11/114029

Klimenko, S., et al. 2016, Phys. Rev. D, 93, 042004, doi: 10.1103/PhysRevD.93.042004

Koehn, H., et al. 2025, Phys. Rev. X, 15, 021014, doi: 10.1103/PhysRevX.15.021014

Koloniari, A. E., Koursoumpa, E. C., Nousi, P., et al. 2025, Mach. Learn. Sci. Tech., 6, 015054, doi: 10.1088/2632-2153/adb5ed
Kremer, K., Spera, M., Becker, D., et al. 2020, Astrophys. J., 903, 45, doi: 10.3847/1538-4357/abb945
Kulkarni, S. F., McMillan, S., & Hut, P. 1993, Nature, 364, 421, doi: 10.1038/364421a0
Kumar, P., & Dent, T. 2024, Phys. Rev. D, 110, 043036, doi: 10.1103/PhysRevD.110.043036
Lalleman, M., Turbang, K., Callister, T., & van Remortel, N. 2025, Astron. Astrophys., 698, A85, doi: 10.1051/0004-6361/202553941
Landry, P., Essick, R., & Chatziioannou, K. 2020, Phys. Rev. D, 101, 123007, doi: 10.1103/PhysRevD.101.123007
Landry, P., & Read, J. S. 2021, Astrophys. J. Lett., 921, L25, doi: 10.3847/2041-8213/ac2f3e
Legred, I., Chatziioannou, K., Essick, R., Han, S., & Landry, P. 2021, Phys. Rev. D, 104, 063003, doi: 10.1103/PhysRevD.104.063003
Leyde, K., Green, S. R., Toubiana, A., & Gair, J. 2024, Phys. Rev. D, 109, 064056, doi: 10.1103/PhysRevD.109.064056
Li, Y.-J., Wang, Y.-Z., Tang, S.-P., Chen, T., & Fan, Y.-Z. 2025, Astrophys. J., 987, 65, doi: 10.3847/1538-4357/add535
Li, Y.-J., Wang, Y.-Z., Tang, S.-P., & Fan, Y.-Z. 2024, Phys. Rev. Lett., 133, 051401, doi: 10.1103/PhysRevLett.133.051401
Li, Y.-J., Wang, Y.-Z., Tang, S.-P., et al. 2022, Astrophys. J. Lett., 933, L14, doi: 10.3847/2041-8213/ac78dd
LIGO Scientific Collaboration, VIRGO Collaboration, & KAGRA Collaboration. 2025, GWTC-4.0: Population Properties of Merging Compact Binaries, Zenodo, doi: 10.5281/zenodo.16911563
Liu, B., & Lai, D. 2021, Mon. Not. Roy. Astron. Soc., 502, 2049, doi: 10.1093/mnras/stab178
Loredo, T. J. 2004, AIP Conf. Proc., 735, 195, doi: 10.1063/1.1835214
Lorenzo-Medina, A., & Dent, T. 2025, Class. Quant. Grav., 42, 045008, doi: 10.1088/1361-6382/ad9c0e
Lousto, C. O., Campanelli, M., Zlochower, Y., & Nakano, H. 2010, Class. Quant. Grav., 27, 114006, doi: 10.1088/0264-9381/27/11/114006
Ma, L., & Fuller, J. 2019, Mon. Not. Roy. Astron. Soc., 488, 4338, doi: 10.1093/mnras/stz2009
—. 2023, Astrophys. J., 952, 53, doi: 10.3847/1538-4357/acdb74
Madau, P., & Dickinson, M. 2014, Ann. Rev. Astron. Astrophys., 52, 415, doi: 10.1146/annurev-astro-081811-125615
Magaña Hernandez, I., & Palmese, A. 2025, Phys. Rev. D, 111, 083031, doi: 10.1103/PhysRevD.111.083031
Mahapatra, P., Chattopadhyay, D., Gupta, A., et al. 2024, Astrophys. J., 975, 117, doi: 10.3847/1538-4357/ad781b
—. 2025a, Phys. Rev. D, 111, 123030, doi: 10.1103/c9l3-gw6w
—. 2025b, Phys. Rev. D, 111, 023013, doi: 10.1103/PhysRevD.111.023013
Mahapatra, P., Gupta, A., Favata, M., Arun, K. G., & Sathyaprakash, B. S. 2021, Astrophys. J. Lett., 918, L31, doi: 10.3847/2041-8213/ac20db
Mali, U., & Essick, R. 2025, Astrophys. J., 980, 85, doi: 10.3847/1538-4357/ad9de7
Mancarella, M., & Gerosa, D. 2025, Phys. Rev. D, 111, 103012, doi: 10.1103/PhysRevD.111.103012
Mandel, I., & Broekgaarden, F. S. 2022, Living Rev. Rel., 25, 1, doi: 10.1007/s41114-021-00034-3
Mandel, I., & de Mink, S. E. 2016, Mon. Not. Roy. Astron. Soc., 458, 2634, doi: 10.1093/mnras/stw379
Mandel, I., & Farmer, A. 2022, Phys. Rept., 955, 1, doi: 10.1016/j.physrep.2022.01.003
Mandel, I., Farr, W. M., Colonna, A., et al. 2017, Mon. Not. Roy. Astron. Soc., 465, 3254, doi: 10.1093/mnras/stw2883
Mandel, I., Farr, W. M., & Gair, J. R. 2019, Mon. Not. Roy. Astron. Soc., 486, 1086, doi: 10.1093/mnras/stz896
Mandel, I., & Fragos, T. 2020, Astrophys. J. Lett., 895, L28, doi: 10.3847/2041-8213/ab8e41
Mandel, I., & Müller, B. 2020, Mon. Not. Roy. Astron. Soc., 499, 3214, doi: 10.1093/mnras/staa3043
Mandel, I., Müller, B., Riley, J., et al. 2020, Mon. Not. Roy. Astron. Soc., 500, 1380, doi: 10.1093/mnras/staa3390
Mapelli, M. 2020, Proc. Int. Sch. Phys. Fermi, 200, 87, doi: 10.3254/ENFI200005
Mapelli, M., Bouffanais, Y., Santoliquido, F., Sedda, M. A., & Artale, M. C. 2022, Mon. Not. Roy. Astron. Soc., 511, 5797, doi: 10.1093/mnras/stac422
Mapelli, M., Santoliquido, F., Bouffanais, Y., et al. 2021, Symmetry, 13, 1678, doi: 10.3390/sym13091678
Mapelli, M., Spera, M., Montanari, E., et al. 2020, Astrophys. J., 888 , 76, doi: 10.3847/1538-4357/ab584d
Marchant, P., Langer, N., Podsiadlowski, P., Tauris, T. M., & Moriya, T. J. 2016, Astron. Astrophys., 588, A50, doi: 10.1051/0004-6361/201628133
Marchant, P., & Moriya, T. 2020, Astron. Astrophys., 640, L18, doi: 10.1051/0004-6361/202038902
Marchant, P., Renzo, M., Farmer, R., et al. 2018, doi: 10.3847/1538-4357/ab3426
Margalit, B., & Metzger, B. D. 2017, Astrophys. J. Lett., 850, L19, doi: 10.3847/2041-8213/aa991c

Martinez, M. A. S., et al. 2020, Astrophys. J., 903, 67, doi: 10.3847/1538-4357/abba25
McKernan, B., Ford, K. E. S., Callister, T., et al. 2022, Mon. Not. Roy. Astron. Soc., 514, 3886, doi: 10.1093/mnras/stac1570
McKernan, B., Ford, K. E. S., Lyra, W., & Perets, H. B. 2012, Mon. Not. Roy. Astron. Soc., 425, 460, doi: 10.1111/j.1365-2966.2012.21486.x
McKernan, B., Ford, K. E. S., O'Shaughnessy, R., & Wysocki, D. 2020, Mon. Not. Roy. Astron. Soc., 494, 1203, doi: 10.1093/mnras/staa740
Mckernan, B., et al. 2018, Astrophys. J., 866, 66, doi: 10.3847/1538-4357/aadae5
Mehta, A. K., Olsen, S., Wadekar, D., et al. 2025, Phys. Rev. D, 111, 024049, doi: 10.1103/PhysRevD.111.024049
Messick, C., et al. 2017, Phys. Rev. D, 95, 042001, doi: 10.1103/PhysRevD.95.042001
Miller, S., Callister, T. A., & Farr, W. 2020, Astrophys. J., 895, 128, doi: 10.3847/1538-4357/ab80c0
Mishra, T., Bhaumik, S., Gayathri, V., et al. 2025, Phys. Rev. D, 111, 023054, doi: 10.1103/PhysRevD.111.023054
Mishra, T., O'Brien, B., Gayathri, V., et al. 2021, Phys. Rev. D, 104, 023014, doi: 10.1103/PhysRevD.104.023014
Mishra, T., et al. 2022, Phys. Rev. D, 105, 083018, doi: 10.1103/PhysRevD.105.083018
Mould, M., & Gerosa, D. 2022, Phys. Rev. D, 105, 024076, doi: 10.1103/PhysRevD.105.024076
Mould, M., Gerosa, D., Broekgaarden, F. S., & Steinle, N. 2022a, Mon. Not. Roy. Astron. Soc., 517, 2738, doi: 10.1093/mnras/stac2859
Mould, M., Gerosa, D., & Taylor, S. R. 2022b, Phys. Rev. D, 106, 103013, doi: 10.1103/PhysRevD.106.103013
Mould, M., Moore, C. J., & Gerosa, D. 2024, Phys. Rev. D, 109, 063013, doi: 10.1103/PhysRevD.109.063013
Nathanail, A., Most, E. R., & Rezzolla, L. 2021, Astrophys. J. Lett., 908, L28, doi: 10.3847/2041-8213/abdfc6
Neijssel, C. J., Vigna-Gómez, A., Stevenson, S., et al. 2019, Mon. Not. Roy. Astron. Soc., 490, 3740, doi: 10.1093/mnras/stz2840
Nitz, A. H., Capano, C., Nielsen, A. B., et al. 2019, Astrophys. J., 872, 195, doi: 10.3847/1538-4357/ab0108
Nitz, A. H., Capano, C. D., Kumar, S., et al. 2021, Astrophys. J., 922, 76, doi: 10.3847/1538-4357/ac1c03
Nitz, A. H., Dal Canton, T., Davis, D., & Reyes, S. 2018, Phys. Rev. D, 98, 024050, doi: 10.1103/PhysRevD.98.024050
Nitz, A. H., Dent, T., Dal Canton, T., Fairhurst, S., & Brown, D. A. 2017, Astrophys. J., 849, 118, doi: 10.3847/1538-4357/aa8f50
Nitz, A. H., Kumar, S., Wang, Y.-F., et al. 2023, Astrophys. J., 946, 59, doi: 10.3847/1538-4357/aca591
Nitz, A. H., Dent, T., Davies, G. S., et al. 2020, Astrophys. J., 891, 123, doi: 10.3847/1538-4357/ab733f
Olejak, A., Fryer, C. L., Belczynski, K., & Baibhav, V. 2022, Mon. Not. Roy. Astron. Soc., 516, 2252, doi: 10.1093/mnras/stac2359
Olejak, A., Klencki, J., Xu, X.-T., et al. 2024, Astron. Astrophys., 689, A305, doi: 10.1051/0004-6361/202450480
Olsen, S., Venumadhav, T., Mushkin, J., et al. 2022, Phys. Rev. D, 106, 043009, doi: 10.1103/PhysRevD.106.043009
Özel, F., & Freire, P. 2016, Ann. Rev. Astron. Astrophys., 54, 401, doi: 10.1146/annurev-astro-081915-023322
Ozel, F., Psaltis, D., Narayan, R., & McClintock, J. E. 2010, Astrophys. J., 725, 1918, doi: 10.1088/0004-637X/725/2/1918
Packet, W. 1981, A&A, 102, 17
Payne, E., Kremer, K., & Zevin, M. 2024, Astrophys. J. Lett., 966, L16, doi: 10.3847/2041-8213/ad3e82
Payne, E., & Thrane, E. 2023, Phys. Rev. Res., 5, 023013, doi: 10.1103/PhysRevResearch.5.023013
Phan, D., Pradhan, N., & Jankowiak, M. 2019. https://arxiv.org/abs/1912.11554
Pierra, G., Mastrogiovanni, S., & Perriès, S. 2024, Astron. Astrophys., 692, A80, doi: 10.1051/0004-6361/202452545
Pompili, L., et al. 2023, Phys. Rev. D, 108, 124035, doi: 10.1103/PhysRevD.108.124035
Portegies Zwart, S. F., & McMillan, S. 2000, Astrophys. J. Lett., 528, L17, doi: 10.1086/312422
Portegies Zwart, S. F., & Yungelson, L. R. 1998, Astron. Astrophys., 332, 173. https://arxiv.org/abs/astro-ph/9710347
Postnov, K. A., & Yungelson, L. R. 2014, Living Rev. Rel., 17, 3, doi: 10.12942/lrr-2014-3
Pratten, G., et al. 2021, Phys. Rev. D, 103, 104056, doi: 10.1103/PhysRevD.103.104056
Qin, Y., Fragos, T., Meynet, G., et al. 2018, Astron. Astrophys., 616, A28, doi: 10.1051/0004-6361/201832839
Qin, Y., Marchant, P., Fragos, T., Meynet, G., & Kalogera, V. 2019, Astrophys. J. Lett., 870, L18, doi: 10.3847/2041-8213/aaf97b
Racine, E. 2008, Phys. Rev. D, 78, 044021, doi: 10.1103/PhysRevD.78.044021
Ramos-Buades, A., Buonanno, A., Estellés, H., et al. 2023, Phys. Rev. D, 108, 124037, doi: 10.1103/PhysRevD.108.124037
Ray, A., Magaña Hernandez, I., Breivik, K., & Creighton, J. 2024. https://arxiv.org/abs/2404.03166

Ray, A., Magaña Hernandez, I., Mohite, S., Creighton, J., & Kapadia, S. 2023a, Astrophys. J., 957, 37, doi: 10.3847/1538-4357/acf452

Ray, A., et al. 2023b. https://arxiv.org/abs/2306.07190

Renzo, M., Farmer, R. J., Justham, S., et al. 2020, Mon. Not. Roy. Astron. Soc., 493, 4333, doi: 10.1093/mnras/staa549

Rezzolla, L., Most, E. R., & Weih, L. R. 2018, Astrophys. J. Lett., 852, L25, doi: 10.3847/2041-8213/aaa401

Riley, J., Mandel, I., Marchant, P., et al. 2021, Mon. Not. Roy. Astron. Soc., 505, 663, doi: 10.1093/mnras/stab1291

Rinaldi, S., & Del Pozzo, W. 2021, Mon. Not. Roy. Astron. Soc., 509, 5454, doi: 10.1093/mnras/stab3224

Rinaldi, S., Del Pozzo, W., Mapelli, M., Lorenzo-Medina, A., & Dent, T. 2024, Astron. Astrophys., 684, A204, doi: 10.1051/0004-6361/202348161

Rinaldi, S., Liang, Y., Demasi, G., Mapelli, M., & Del Pozzo, W. 2025. https://arxiv.org/abs/2506.05929

Rodriguez, C. L., Amaro-Seoane, P., Chatterjee, S., & Rasio, F. A. 2018, Phys. Rev. Lett., 120, 151101, doi: 10.1103/PhysRevLett.120.151101

Rodriguez, C. L., Chatterjee, S., & Rasio, F. A. 2016a, Phys. Rev. D, 93, 084029, doi: 10.1103/PhysRevD.93.084029

Rodriguez, C. L., & Loeb, A. 2018, Astrophys. J. Lett., 866, L5, doi: 10.3847/2041-8213/aae377

Rodriguez, C. L., Morscher, M., Pattabiraman, B., et al. 2015, Phys. Rev. Lett., 115, 051101, doi: 10.1103/PhysRevLett.115.051101

Rodriguez, C. L., Zevin, M., Amaro-Seoane, P., et al. 2019, Phys. Rev. D, 100, 043027, doi: 10.1103/PhysRevD.100.043027

Rodriguez, C. L., Zevin, M., Pankow, C., Kalogera, V., & Rasio, F. A. 2016b, Astrophys. J. Lett., 832, L2, doi: 10.3847/2041-8205/832/1/L2

Romero-Shaw, I. M., Kremer, K., Lasky, P. D., Thrane, E., & Samsing, J. 2021, Mon. Not. Roy. Astron. Soc., 506, 2362, doi: 10.1093/mnras/stab1815

Romero-Shaw, I. M., Thrane, E., & Lasky, P. D. 2022, Publ. Astron. Soc. Austral., 39, e025, doi: 10.1017/pasa.2022.24

Roulet, J., & Zaldarriaga, M. 2019, Mon. Not. Roy. Astron. Soc., 484, 4216, doi: 10.1093/mnras/stz226

Roy, S. K., van Son, L. A. C., & Farr, W. M. 2025. https://arxiv.org/abs/2507.01086

Ruiz, M., Shapiro, S. L., & Tsokaros, A. 2018, Phys. Rev. D, 97, 021501, doi: 10.1103/PhysRevD.97.021501

Rutherford, N., et al. 2024, Astrophys. J. Lett., 971, L19, doi: 10.3847/2041-8213/ad5f02

Sachdev, S., et al. 2019. https://arxiv.org/abs/1901.08580

Sadiq, J., Dent, T., & Gieles, M. 2024, Astrophys. J., 960, 65, doi: 10.3847/1538-4357/ad0ce6

Sadiq, J., Dent, T., & Lorenzo-Medina, A. 2025a. https://arxiv.org/abs/2502.06451

—. 2025b. https://arxiv.org/abs/2506.02250

Sadiq, J., Dent, T., & Wysocki, D. 2022, Phys. Rev. D, 105, 123014, doi: 10.1103/PhysRevD.105.123014

Safarzadeh, M., Farr, W. M., & Ramirez-Ruiz, E. 2020, Astrophys. J., 894, 129, doi: 10.3847/1538-4357/ab80be

Sakon, S., et al. 2024, Phys. Rev. D, 109, 044066, doi: 10.1103/PhysRevD.109.044066

Salpeter, E. E. 1955, Astrophys. J., 121, 161, doi: 10.1086/145971

Santini, A., Gerosa, D., Cotesta, R., & Berti, E. 2023, Phys. Rev. D, 108, 083033, doi: 10.1103/PhysRevD.108.083033

Schiebelbein-Zwack, A., & Fishbach, M. 2024, Astrophys. J., 970, 128, doi: 10.3847/1538-4357/ad5353

Schmidt, P., Hannam, M., & Husa, S. 2012, Phys. Rev. D, 86, 104063, doi: 10.1103/PhysRevD.86.104063

Schmidt, P., Hannam, M., Husa, S., & Ajith, P. 2011, Phys. Rev. D, 84, 024046, doi: 10.1103/PhysRevD.84.024046

Schmidt, P., Ohme, F., & Hannam, M. 2015, Phys. Rev. D, 91, 024043, doi: 10.1103/PhysRevD.91.024043

Shen, Y., et al. 2023, Astrophys. J., 945, 41, doi: 10.3847/1538-4357/acb7de

Sigurdsson, S., & Hernquist, L. 1993, Nature, 364, 423, doi: 10.1038/364423a0

Speagle, J. S. 2020, Mon. Not. Roy. Astron. Soc., 493, 3132, doi: 10.1093/mnras/staa278

Stegmann, J., & Antonini, F. 2021, Phys. Rev. D, 103, 063007, doi: 10.1103/PhysRevD.103.063007

Steinle, N., & Kesden, M. 2021, Phys. Rev. D, 103, 063032, doi: 10.1103/PhysRevD.103.063032

Stevenson, S. 2022, Astrophys. J. Lett., 926, L32, doi: 10.3847/2041-8213/ac5252

Stone, N. C., Metzger, B. D., & Haiman, Z. 2017, Mon. Not. Roy. Astron. Soc., 464, 946, doi: 10.1093/mnras/stw2260

Sukhbold, T., Ertl, T., Woosley, S. E., Brown, J. M., & Janka, H. T. 2016, Astrophys. J., 821, 38, doi: 10.3847/0004-637X/821/1/38

Szemraj, L., & Biscoveanu, S. 2025. https://arxiv.org/abs/2507.23663

Talbot, C., Farah, A., Galaudage, S., Golomb, J., & Tong, H. 2025a, J. Open Source Softw., 10, 7753, doi: 10.21105/joss.07753

Talbot, C., & Golomb, J. 2023, Mon. Not. Roy. Astron. Soc., 526, 3495, doi: 10.1093/mnras/stad2968

Talbot, C., & Thrane, E. 2017, Phys. Rev. D, 96, 023012, doi: 10.1103/PhysRevD.96.023012

—. 2022, Astrophys. J., 927, 76, doi: 10.3847/1538-4357/ac4bc0
Talbot, C., et al. 2025b. https://arxiv.org/abs/2508.11091
Tauris, T. M. 2022, Astrophys. J., 938, 66, doi: 10.3847/1538-4357/ac86c8
Thomas, L. M., Schmidt, P., & Pratten, G. 2021, Phys. Rev. D, 103, 083022, doi: 10.1103/PhysRevD.103.083022
Thompson, T. A., et al. 2018, doi: 10.1126/science.aau4005
Thrane, E., & Talbot, C. 2019, Publ. Astron. Soc. Austral., 36, e010, doi: 10.1017/pasa.2019.2
Tiwari, V. 2018, Class. Quant. Grav., 35, 145009, doi: 10.1088/1361-6382/aac89d
—. 2021, Class. Quant. Grav., 38, 155007, doi: 10.1088/1361-6382/ac0b54
—. 2022, Astrophys. J., 928, 155, doi: 10.3847/1538-4357/ac589a
—. 2023, Mon. Not. Roy. Astron. Soc., 527, 298, doi: 10.1093/mnras/stad3155
Tiwari, V., & Fairhurst, S. 2021, Astrophys. J. Lett., 913, L19, doi: 10.3847/2041-8213/abfbe7
Tiwari, V., Klimenko, S., Necula, V., & Mitselmakher, G. 2016, Class. Quant. Grav., 33, 01LT01, doi: 10.1088/0264-9381/33/1/01LT01
Tong, H., Galaudage, S., & Thrane, E. 2022, Phys. Rev. D, 106, 103019, doi: 10.1103/PhysRevD.106.103019
Torniamenti, S., Mapelli, M., Périgois, C., et al. 2024, Astron. Astrophys., 688, A148, doi: 10.1051/0004-6361/202449272
Toubiana, A., Katz, M. L., & Gair, J. R. 2023, Mon. Not. Roy. Astron. Soc., 524, 5844, doi: 10.1093/mnras/stad2215
Tout, C. A., & Pringle, J. E. 1992, MNRAS, 256, 269, doi: 10.1093/mnras/256.2.269
Trani, A. A., Tanikawa, A., Fujii, M. S., Leigh, N. W. C., & Kumamoto, J. 2021, Mon. Not. Roy. Astron. Soc., 504, 910, doi: 10.1093/mnras/stab967
Tsukada, L., et al. 2023, Phys. Rev. D, 108, 043004, doi: 10.1103/PhysRevD.108.043004
Turbang, K., Lalleman, M., Callister, T. A., & van Remortel, N. 2024, Astrophys. J., 967, 142, doi: 10.3847/1538-4357/ad3d5c
Usman, S. A., et al. 2016, Class. Quant. Grav., 33, 215004, doi: 10.1088/0264-9381/33/21/215004
van den Heuvel, E. P. J., Portegies Zwart, S. F., & de Mink, S. E. 2017, Mon. Not. Roy. Astron. Soc., 471, 4256, doi: 10.1093/mnras/stx1430
van Son, L. A. C., de Mink, S. E., Chruslinska, M., et al. 2022a, doi: 10.3847/1538-4357/acbf51
van Son, L. A. C., de Mink, S. E., Broekgaarden, F. S., et al. 2020, Astrophys. J., 897, 100, doi: 10.3847/1538-4357/ab9809
van Son, L. A. C., de Mink, S. E., Callister, T., et al. 2022b, Astrophys. J., 931, 17, doi: 10.3847/1538-4357/ac64a3
van Son, L. A. C., de Mink, S. E., Renzo, M., et al. 2022c, Astrophys. J., 940, 184, doi: 10.3847/1538-4357/ac9b0a
Varma, V., Field, S. E., Scheel, M. A., et al. 2019, Phys. Rev. Research., 1, 033015, doi: 10.1103/PhysRevResearch.1.033015
Venumadhav, T., Zackay, B., Roulet, J., Dai, L., & Zaldarriaga, M. 2019, Phys. Rev. D, 100, 023011, doi: 10.1103/PhysRevD.100.023011
—. 2020, Phys. Rev. D, 101, 083030, doi: 10.1103/PhysRevD.101.083030
Vijaykumar, A., Fishbach, M., Adhikari, S., & Holz, D. E. 2024, Astrophys. J., 972, 157, doi: 10.3847/1538-4357/ad6140
Vitale, S., Biscoveanu, S., & Talbot, C. 2022, Astron. Astrophys., 668, L2, doi: 10.1051/0004-6361/202245084
Vitale, S., Farr, W. M., Ng, K., & Rodriguez, C. L. 2019, Astrophys. J. Lett., 886, L1, doi: 10.3847/2041-8213/ab50c0
Vitale, S., Gerosa, D., Farr, W. M., & Taylor, S. R. 2020, doi: 10.1007/978-981-15-4702-7_45-1
Vitale, S., Lynch, R., Sturani, R., & Graff, P. 2017, Class. Quant. Grav., 34, 03LT01, doi: 10.1088/1361-6382/aa552e
Vitale, S., & Mould, M. 2025. https://arxiv.org/abs/2505.14875
Wadekar, D., Roulet, J., Venumadhav, T., et al. 2023. https://arxiv.org/abs/2312.06631
Wang, Y.-H., McKernan, B., Ford, S., et al. 2021, Astrophys. J. Lett., 923, L23, doi: 10.3847/2041-8213/ac400a
Wang, Y.-Z., Li, Y.-J., Vink, J. S., et al. 2022, Astrophys. J. Lett., 941, L39, doi: 10.3847/2041-8213/aca89f
Wiktorowicz, G., Wyrzykowski, Ł., Chruslinska, M., et al. 2019, doi: 10.3847/1538-4357/ab45e6
Winch, E. R. J., Vink, J. S., Higgins, E. R., & Sabhahitf, G. N. 2024, Mon. Not. Roy. Astron. Soc., 529, 2980, doi: 10.1093/mnras/stae393
Woosley, S. E. 2017, Astrophys. J., 836, 244, doi: 10.3847/1538-4357/836/2/244
Woosley, S. E., & Heger, A. 2021, Astrophys. J. Lett., 912, L31, doi: 10.3847/2041-8213/abf2c4
Wysocki, D., Gerosa, D., O'Shaughnessy, R., et al. 2018, Phys. Rev. D, 97, 043014, doi: 10.1103/PhysRevD.97.043014

Wysocki, D., Lange, J., & O'Shaughnessy, R. 2019, Phys. Rev. D, 100, 043012, doi: 10.1103/PhysRevD.100.043012

Yang, Y., Bartos, I., Haiman, Z., et al. 2019, Astrophys. J., 876, 122, doi: 10.3847/1538-4357/ab16e3

Ye, C. S., & Fishbach, M. 2024, Astrophys. J., 967, 62, doi: 10.3847/1538-4357/ad3ba8

Ye, C. S., Fong, W.-f., Kremer, K., et al. 2020, Astrophys. J. Lett., 888, L10, doi: 10.3847/2041-8213/ab5dc5

Ye, C. S., Kremer, K., Ransom, S. M., & Rasio, F. A. 2024, Astrophys. J., 975, 77, doi: 10.3847/1538-4357/ad76a0

Zackay, B., Venumadhav, T., Dai, L., Roulet, J., & Zaldarriaga, M. 2019, Phys. Rev. D, 100, 023007, doi: 10.1103/PhysRevD.100.023007

Zaldarriaga, M., Kushnir, D., & Kollmeier, J. A. 2018, Mon. Not. Roy. Astron. Soc., 473, 4174, doi: 10.1093/mnras/stx2577

Zevin, M., & Bavera, S. S. 2022, Astrophys. J., 933, 86, doi: 10.3847/1538-4357/ac6f5d

Zevin, M., Bavera, S. S., Berry, C. P. L., et al. 2021, Astrophys. J., 910, 152, doi: 10.3847/1538-4357/abe40e

Zhang, R. C., Fragione, G., Kimball, C., & Kalogera, V. 2023, Astrophys. J., 954, 23, doi: 10.3847/1538-4357/ace4c1

Ziosi, B. M., Mapelli, M., Branchesi, M., & Tormen, G. 2014, Mon. Not. Roy. Astron. Soc., 441, 3703, doi: 10.1093/mnras/stu824